arXiv:1607.08876v2 [math.AG] 8 Aug 2016

The noncommutative geometry of elliptic difference equations

Eric M. Rains
Department of Mathematics, California Institute of Technology

August 7, 2016

Abstract

We give a new construction of noncommutative rational surfaces via elliptic difference operators, which enables us to attach a 1-parameter noncommutative deformation to any projective rational surface with smooth anticanonical curve. The result turns out to agree with a construction implicit in work of Van den Bergh (as blowups of noncommutative Hirzebruch surfaces), but the construction, by giving an alternate interpretation of morphisms between line bundles, enables one to prove a number of new facts about these surfaces. In particular, we show that they are noncommutative smooth proper surfaces in the sense of Chan and Nyman, with projective Quot schemes, that moduli spaces of simple sheaves are Poisson and that moduli spaces classifying semistable sheaves of rank 0 or 1 are projective. We further show that the action of $\mathrm{SL}_2(\mathbb{Z})$ as derived autoequivalences of rational elliptic surfaces extends to an action as derived equivalences of surfaces in our family with $K^2 = 0$.

This construction also enables a number of new results in the theory of special functions. Since morphisms between line bundles on these noncommutative surfaces are difference operators, this leads to an interpretation of certain moduli spaces of sheaves on the surfaces as moduli spaces of difference equations. When the moduli space is a single point, the equation is rigid, and we show that one can construct solutions as generalized hypergeometric integrals. More generally, twisting by line bundles turns out to give rise to isomonodromy deformations, and thus such moduli spaces naturally produce Lax pairs. In particular, we show that when the moduli space is 2-dimensional, it gives rise to a Lax pair for the elliptic Painlevé equation; using the derived equivalences, this gives such a Lax pair for every rational number, of order twice the denominator. We also construct an elliptic analogue of the Riemann-Hilbert correspondence, an analytic equivalence between categories of sheaves corresponding to elliptic difference equations, swapping the role of the shift of the equation and the nome of the curve.

Contents

1 Introduction **2**

2 Difference equations as modules **11**

3 Flatness (even Hirzebruch case) **17**

4 Torsion modules and sheaves I **25**

5 F_1 and $\mathbb{P}^2$ **38**

6 Surfaces with $K_X^2 = 7$ **46**

7 General blowups **56**

8 Surjectivity of composition **65**

9 Torsion modules and sheaves II **72**

10 Cohomology and duality **90**

11 Moduli spaces **105**
11.1 Poisson structures 105
11.2 Hilbert schemes 120
11.3 Semistable 1-dimensional sheaves 130
11.4 Lax pairs for elliptic Painlevé 137

12 Derived equivalences **140**

13 Sheaves as difference equations **155**
13.1 Spaces of meromorphic solutions 155
13.2 Sheaves of holomorphic solutions 162
13.3 The elliptic Riemann-Hilbert correspondence 172
13.4 The Fourier transform 179
13.5 Generalized hypergeometric functions 187

Appendix A An integrable system on relations **189**

Appendix B Blowing up $\mathbb{P}^n$ **194**

1 Introduction

A particularly striking characterization of Gauss' hypergeometric function ${}_2F_1$ is as the solution of a second-order Fuchsian differential equation with three singular points. In fact, such equations are *rigid*, in the sense that if one specifies the local behavior of the equation near the singular points, this is enough to determine the equation up to isomorphism. If one adds an additional singular point, one no longer has rigidity, but now acquires a new structure in the form of a flow (the Painlevé VI equation) that moves the fourth singular point without changing the monodromy of the equation. Both of these structures generalize considerably; e.g., the ${}_nF_{n-1}$ functions satisfy rigid n-th order equations, and there are flows analogous to Painlevé on higher-dimensional moduli spaces of equations. Moreover, there are discrete analogues: the contiguity relations for ${}_nF_{n-1}$ can be viewed as rigid *difference* equations, and similar things hold for the q-hypergeometric ${}_n\phi_{n-1}$; there are also discrete and q-analogues of Painlevé constructed as isomorphisms between moduli spaces of difference or q-difference equations.

We thus see that much of the modern theory of special functions is closely related to moduli spaces of differential or difference equations. In particular, we might hope that a better understanding of these moduli spaces would lead to new insights into existing special functions, as well as to new special functions and integrable systems. Indeed, both the hypergeometric equations and the Painlevé equations actually extend beyond the q-difference case to the *elliptic* difference case, in which multiplication by q is replaced by translation on an elliptic curve. However, Sakai's construction of the elliptic Painlevé equation in [54] was based on the geometry of rational surfaces rather than on difference equations; though later work produced such interpretations [67, 50], it has been unclear how to extend this to analogues of the Garnier equations (flows on second-order Fuchsian equations with $n > 4$ singular points) or other generalizations.

In [47], the author gave a first pass at understanding such moduli spaces, by giving a correspondence between various sorts of difference and differential equations and sheaves on Hirzebruch surfaces. However, though this behaves well on an open subset of the moduli space of equations, it leads to an incorrect compactification. For instance, one can view a differential equation as a connection on a vector bundle, and the correspondence in [47] behaves well only when the vector bundle is trivial. Now, differential equations with rational coefficients can be interpreted as D-modules on $\mathbb{P}^1$, which we can equivalently view as sheaves on a noncommutative deformation of $\mathbb{A}^1 \times \mathbb{P}^1$. This suggests that to fix this issue, we should replace the Hirzebruch surfaces with noncommutative deformations. (This is further supported by the construction in [39] of the elliptic Painlevé equation via sheaves on a noncommutative projective plane.)

The natural hope would be that the homogeneous coordinate ring of such a noncommutative Hirzebruch surface would have a faithful representation as an algebra of difference operators. This, however, encounters a significant difficulty. Although many early constructions of noncommutative schemes were obtained from flat deformations of graded algebras, there is in general *no* graded algebra associated to a given noncommutative scheme. The difficulty here is that although line bundles can often be extended along a noncommutative deformation, they deform only as left modules rather than as bimodules. In practice, only multiples of the canonical bundle extend as bimodules[1], and thus we only obtain a homogeneous coordinate ring when some such multiple is very ample. Thus in particular, we would only expect to have a homogeneous coordinate ring associated to a deformation of the Hirzebruch surfaces $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$ and F_1, as these are the only Fano Hirzebruch surfaces (i.e., such that the anticanonical bundle is ample). This issue becomes even worse when we realize that [47] tells us that the most natural way to encode the singularity data is to blow up our surface and consider (1-dimensional) sheaves which are *disjoint* from the anticanonical curve, since the existence of such sheaves ensures that the anticanonical divisor cannot possibly be ample!

The standard way [11] to resolve this issue is to suppose that there is some ample bundle $\mathcal{O}(1)$ such that all powers $\mathcal{O}(k)$ extend to the noncommutative deformation. In that case, we can construct not an algebra but a category, in which the objects are integers, and the Hom spaces are defined by

$$\mathrm{Hom}(k, l) = \mathrm{Hom}(\mathcal{O}(-l), \mathcal{O}(-k)). \tag{1.1}$$

Such a category with objects indexed by $\mathbb{Z}$ and all morphisms having nonnegative degree is called a $\mathbb{Z}$-algebra, and one can indeed recover the (category of sheaves on the) original noncommutative scheme as a quotient of the category of modules over this $\mathbb{Z}$-algebra. Of course, in addition to the choice of original ample bundle and the requirement that the bundles extend, there is also in principle a choice of such extension for each l. These are both serious issues in general, though in our case it is reasonable to hope that neither issue arises. Indeed, it follows from the deformation theory of abelian categories [33] that the infinitesimal obstructions to extending a line bundle lie in $H^2(\mathcal{O}_X)$, while the choice of extension is controlled by $H^1(\mathcal{O}_X)$. In our case, since we are deforming a rational surface, both cohomology groups vanish, so neither problem arises for a formal deformation. (There is still a mild nonuniqueness, in that we can multiply each line bundle by a unit, but this freedom is easy enough to deal with.)

Although this approach indeed works in our case, the requirement to choose an ample bundle can make it difficult to detect when two $\mathbb{Z}$-algebras induce isomorphic schemes (indeed, it can be difficult to determine when two *commutative* graded algebras correspond to isomorphic projective schemes). As a result, it will be convenient to put off making such a choice as long as possible, suggesting that we should deform *all* ample line bundles. In fact, since we cannot hope for the

[1] In the cases we study here, this follows from Proposition 12.22 below.

ample cone to remain constant under deformation, the simplest thing to do is simply deform all line bundles, ample or otherwise. Given a smooth projective rational surface X, the Picard group $\mathrm{Pic}(X) \cong \mathbb{Z}^m$ for some m (to be precise, $m = 10 - K_X^2$). If for each $\lambda \in \mathrm{Pic}(X)$, we choose a representative $\mathcal{L}_\lambda$ of that isomorphism class of line bundles, then we obtain a category associated to X as follows. The objects of the category are the vectors $\lambda \in \mathrm{Pic}(X)$, while for each λ, μ,

$$\mathrm{Hom}(\lambda, \mu) := \mathrm{Hom}(\mathcal{L}_{-\mu}, \mathcal{L}_{-\lambda}). \tag{1.2}$$

Given a deformation of X such that all line bundles extend uniquely (up to isomorphism), we obtain a corresponding deformation of this category. This introduces some other issues (there are difficulties in defining the corresponding category of sheaves), but has distinct advantages for present purposes. In particular, for a large class of projective rational surfaces X (those having a smooth anticanonical curve), we will be able to construct a flat noncommutative deformation of this category. Moreover, the entire category will (by construction) have a faithful representation in suitable difference operators. (We will extend this in [46] to an analogous family of *multivariate* difference operators, giving a 2-parameter deformation of $\mathrm{Sym}^n(X)$ for any projective rational surface X with a smooth anticanonical curve.)

Note that one cost of working with the full $\mathrm{Pic}(X)$ is that our $\mathbb{Z}^m$-algebras are not (locally) finitely generated in general; indeed, even in the commutative setting the full multihomogeneous coordinate ring needs at least one generator for every -1-curve on the surface, and there can easily be infinitely many such curves once $K^2 \leq 0$. Thus we mostly must eschew the traditional approach the to construction of deformations via generators and relations. Here we are helped greatly by the fact that we are specifying a sub-$\mathbb{Z}^m$-algebra of a $\mathbb{Z}^m$-algebra of difference operators, so that even in the finitely generated case we need not bother with relations, and in general it suffices to specify subspaces respected by composition.

This construction should be contrasted with one implicit in [64, 65]: to construct a noncommutative rational surface, one can simply construct a noncommutative Hirzebruch surface (or noncommutative $\mathbb{P}^2$) and then repeatedly blow it up. This has some advantages over our approach, in that it is clear the resulting categories of sheaves are noetherian and otherwise well-behaved, but compensatory disadvantages, most significantly that it is difficult in this approach to understand when different such constructions result in isomorphic schemes. For instance, the commutative surface $\mathbb{P}^1 \times \mathbb{P}^1$ can be viewed as a Hirzebruch surface in two different ways, but this is difficult to see if we are given the surface *as* a Hirzebruch surface, and only becomes more difficult in the noncommutative setting. In contrast, we will find that our deformations have an explicit symmetry corresponding to this phenomenon, and in general will be isomorphic whenever the commutative surfaces being deformed are isomorphic. We will in fact see below that, at least in the smooth anticanonical curve case, both approaches are actually equivalent: the schemes associated to our noncommutative rational surfaces can indeed be obtained via Van den Bergh's constructions. However, the corresponding representation in difference operators is not only new, but useful in proving symmetries of the constructions.

Though the construction à la Van den Bergh works for any rational surface with a choice of Poisson structure (i.e., an infinitesimal noncommutative deformation), we will be restricting our attention in the present work to the case when the kernel of the Poisson structure is a smooth (anticanonical) curve C. Since we many of our arguments use this curve to reduce to commutative questions, the smooth case is the simplest to deal with. Of course, this excludes cases corresponding to the vast majority of the work in special functions, most of which corresponds to differential equations, for which [47] tells us the anticanonical curve is not even reduced. We will thus return to that case in [44], replacing the present $\mathbb{Z}^m$-algebra approach by one via derived categories based on insights gained below.

Given a smooth genus 1 curve C, the general rational surface X with an embedding $C \to X$ as an anticanonical curve can be constructed in the following way from data on C, specifically line bundles $\eta, \eta' \in \mathrm{Pic}^2(C)$ and points $x_1, \ldots, x_m \in C$. The line bundle η determines a degree 2 map $\rho : C \to \mathbb{P}^1$, and thus $\rho_* \eta'$ is a rank 2 vector bundle on this $\mathbb{P}^1$. The corresponding Hirzebruch surface X_0 contains C in a natural way, and we in fact find that C is an anticanonical curve on X_0. Blowing up the point $x_1 \in C \subset X_0$ gives a surface X_1 which still contains C as an anticanonical curve, so that we may iterate to obtain a surface $X = X_m$. In this way, we obtain every rational surface containing C anticanonically, except for F_1 and $\mathbb{P}^2$, which we may obtain by blowing down suitable -1-curves on X_1. Note that this construction comes with a natural choice of basis for $\mathrm{Pic}(X)$: we have classes f, s, e_1,...,e_m, such that f is the class of a fiber of ρ, s is the relative $\mathcal{O}(1)$, and e_i is the exceptional curve of the i-th blowup. The intersection form in this basis is

$$s \cdot f = 1, \quad e_i^2 = -1, \tag{1.3}$$

with all other intersections 0. There are two other bases of interest. If we start with F_1, we have a similar basis, but now the minimal section is a -1-curve so $s^2 = -1$; starting with $\mathbb{P}^2$ gives a basis $h, e_0, e_1, \dot{,} s, e_m$ with

$$h^2 = 1, \quad e_i^2 = -1. \tag{1.4}$$

Note that since a Poisson structure on X vanishing on C depends on a choice of differential on C, the infinitesimal noncommutative deformations of X are parametrized by the Lie algebra of $\mathrm{Pic}^0(C)$; this suggests that the actual deformations should be parametrized by a point $q \in \mathrm{Pic}^0(C)$. (In this, we may also be guided by the fact that the standard "Sklyanin" noncommutative $\mathbb{P}^2$ is parametrized in precisely that way.)

Thus for each C and integer $m \geq 0$, we wish to construct a category with objects in $\mathbb{Z}^{m+2}$ depending on the above data together with a point $q \in \mathrm{Pic}^0(C)$. As we vary the parameters (including C itself), this deformation should be flat in a suitable sense; we will in fact be able to arrange for *every* Hom space to be a flat sheaf on parameter space. When q becomes trivial, we want to recover the category associated to the original surface X_m. The cost of flatness is that this will only literally be true when X_m has no -2-curves, but any Hom space associated to an *ample* divisor will have the correct limit as q becomes trivial, and this is all we really need. (To be precise, the Hom spaces in the commutative limit are the direct images of corresponding line bundles on the universal surface, shown to be flat in [47]; thus any Hom space corresponding to an acyclic line bundle will have the correct limit, while in general we will obtain only a subspace of the true Hom space. This is the analogue of working with a model for a commutative variety which is not projectively normal.)

As noted in [47], the surface does not uniquely determine the data C, η, η', x_1, ..., x_m. In addition to the obvious freedom to replace C by an isomorphic curve (which will only require an easy functoriality argument on our part), there is also an action of an in general infinite Coxeter group $W(E_{m+1})$ on the data that preserves the isomorphism class of X_m. This group is generated by the following reflections: (1) For each $1 \leq i < m$, we may swap x_i and x_{i+1}, (2) We may replace η', x_1, x_2 by $\eta' + \eta - x_1 - x_2$, $\eta - x_2$, $\eta - x_1$, and (3) we may swap η and η'. Indeed, (1) if $x_i \neq x_{i+1}$, then the locality of blowing up implies that the order of blowing the two points up is irrelevant, (2) if we blow up x_1 and x_2 and blow down the corresponding fibers, we obtain a new, generically isomorphic, Hirzebruch surface (with the given parameters), and (3) if $\eta \neq \eta'$, then $X_0 \cong \mathbb{P}^1 \times \mathbb{P}^1$, and swapping η and η' simply switches to the other interpretation of X_0 as a Hirzebruch surface. (For the somewhat harder fact that these are the only isomorphisms, see [47] or Proposition 12.22 below.) Luckily, our flat deformation turns out to be invariant under the action of this Coxeter group. More precisely, the Coxeter group acts on the basis of $\mathrm{Pic}(X)$ as

well as on the parameters, and the combined action leaves the Hom spaces invariant. The results of [47] suggest that the reflections of type (1) will have no effect beyond bookkeeping, while the reflection of type (2) will essentially be just conjugation by the solution of a certain first-order difference equation (more precisely, gauging by a scalar 1-cocycle). The main difficulty will thus be to understand how the reflections of type (3) act. Luckily, just as in the commutative setting, we will only need to understand the analogue of $\mathbb{P}^1 \times \mathbb{P}^1$ to make this work.

To construct the deformations of $\mathbb{P}^1 \times \mathbb{P}^1$, we can be guided by the fact mentioned above that certain kinds of elliptic difference equations should correspond to sheaves on such deformations; this will indeed give enough guidance to allow us to construct the deformed category in this setting. Moreover, since we have a fair understanding from [47] of how elementary transformations should appear in the difference operator interpretation, we can use this to determine what blowing up a point should do, and thus (with some additional work) how the general $\mathbb{Z}^{m+2}$-algebra in our family should be defined.

At this point, we must now contend with what we gave up when working with $\mathbb{Z}^{m+2}$-algebras. The problem here is that a $\mathbb{Z}^{m+2}$-algebra is the analogue of a multigraded algebra, and even in the commutative setting, there is no natural way to turn a multigraded algebra into a scheme! If one starts with a scheme, then one certainly has a functor from the category of modules over the multihomogeneous coordinate ring into the category of quasicoherent sheaves: restrict to the graded subalgebra corresponding to an ample divisor class, and then translate to a sheaf in the usual way. The issue is thus to determine the kernel of this functor: i.e., which modules over the $\mathbb{Z}^{m+2}$-algebra correspond to the trivial quasicoherent sheaf? The point now is that different choices of ample bundle in principle give rise to different "torsion" subcategories. Luckily, we can give a uniform definition of torsion elements in modules over our $\mathbb{Z}^{m+2}$, and show that this agrees with the notion coming from any choice of ample divisor class. Moreover, one can recover the construction à la Van den Bergh via a particular inductive definition of "torsion", which again agrees with our uniform description.

We thus obtain a flat family of noncommutative schemes (more precisely, the corresponding categories of (quasi)coherent sheaves) associated to our construction. Moreover, by using both what we know from the $\mathbb{Z}^{m+2}$-algebra construction and what Van den Bergh tells us about blowups, we are able to show that these schemes are "noncommutative smooth proper surfaces" in the sense of [14]. In particular, there is an analogue of Serre duality, in which twisting by the canonical bundle is replaced by the application of a certain invertible endofunctor. (Moreover, our surfaces all have an "anticanonical" natural transformation, from this canonical functor to the identity.) We are also able to characterize the analogues of the effective and nef cones, and show that a divisor is ample iff it is in the interior of the nef cone. In addition, we prove a stronger version of the "halal Hilbert scheme" condition: Quot schemes on our surfaces are not merely countable unions of projective schemes, but are in fact projective.

Now that we have well-behaved noncommutative schemes, the next step is to investigate the corresponding moduli spaces of sheaves. In the commutative setting, the presence of an anticanonical curve induces Poisson structures [63, 12, 25] on the moduli spaces, and [45] moreover showed that a number of natural morphisms between such moduli spaces preserved the Poisson structure. Since our surfaces have an anticanonical natural transformation, the construction in [63] of a pairing on the cotangent sheaf carries over directly, and it remains only for us to show that it is alternating and satisfies the Jacobi identity. We in fact find that the moduli space of "simple" coherent sheaves (i.e., with no nontrivial endomorphisms) is indeed Poisson in this way, with symplectic leaves determined by how the sheaves meet the anticanonical curve. We already know from [47] that there are open subsets of the moduli space that are independent of q, and we can show that this not only extends to appropriate torsion-free sheaves, but that the correspondence preserves the Poisson

structure. The key insight there is that the correspondence actually works at the level of derived categories: the kernel of the derived global sections functor is actually independent of q.

One disadvantage of the moduli space of simple sheaves is that it is only an algebraic space. We would thus like to replace it in general by the moduli space of stable sheaves (which would inherit the Poisson structure) or semistable sheaves. Although we are able to show that these are moderately well-behaved (separated, resp. proper) schemes, there are difficulties in showing projectivity in general. There are two special cases in which we can control things, however. The first is the analogue of the Hilbert scheme of points on X. Following [36], this is not simply the Quot scheme of 0-dimensional quotients of $\mathcal{O}_X$, but rather the moduli space of sheaves with the same numerical invariants as an ideal sheaf of n points. We show that in general, this is a smooth, irreducible, rational, projective variety of dimension $2n$, with a natural Poisson structure. Two low-dimensional cases are of particular interest: the $n = 0$ case leads to a numerical characterization of line bundles, while the $n = 1$ case gives us a smooth Poisson rational surface associated to any of our noncommutative rational surfaces. In particular, we find that not only does our family of noncommutative rational surfaces have one more parameter than the family of commutative rational surfaces, but it is in fact in a natural way a $\mathrm{Pic}^0(C)$-bundle over that family. For general n, this construction gives a new deformation of the Hilbert scheme of a rational surface, extending the construction of [36] for $\mathbb{P}^2$.

The other case in which we have sufficient boundedness to control the moduli space is the case of 1-dimensional sheaves; this is particularly important to us since the moduli spaces of difference equations are open subschemes. Here we are again able to prove sufficient boundedness and inequalities on global sections to show that the moduli space of semistable sheaves is projective. If $\chi(M) = 0$, we find as expected from [47] that one component of the moduli space is birational to a commutative moduli space, while if $\chi(M) = \pm 1$, we obtain a birational map between the moduli space of 1-dimensional sheaves and an appropriate Hilbert scheme. One interesting consequence of this is a description of certain special cases of the isomonodromy transformations in terms of arithmetic on moving Jacobians as in [29, 39].

When the moduli space is 2-dimensional and there is a universal family, we can say much more: the moduli space is a rational surface with $K^2 = 0$ and a smooth anticanonical curve. In particular, any such case gives rise to a Lax pair for an elliptic Painlevé equation. We moreover have enough control over the resulting surface to be able to say precisely which surface arises, and thus can describe the full family of Lax pairs arising this way for a given instance of the elliptic Painlevé equation. We find as a result that there is such a Lax pair associated to any rational number (which in a sense includes the case of ∞, corresponding to Sakai's original construction).

The key idea to control the moduli space in the 2-dimensional case is that when the moduli space is 2-dimensional, the universal family gives rise to a derived functor à la Fourier-Mukai, and in fact that derived functor is an equivalence. It turns out that any such derived equivalence respects the canonical functor and the anticanonical natural transformation, and thus respects restriction to the anticanonical curve; this makes it relatively straightforward to read off the parameters. Of course, by composing these derived equivalences with commutative surfaces, we also obtain a number of derived equivalences between noncommutative surfaces. We find in this way that there is an action of $\mathrm{SL}_2(\mathbb{Z})$ on the parameters in such a way that any two surfaces in the same orbit are derived equivalent whenever one of the surfaces in the orbit is commutative.

In fact, it turns out that that last condition is unnecessary: we can construct the desired equivalences using known semiorthogonal decompositions and thus obtain an action of an extension of $\mathrm{SL}_2(\mathbb{Z})$ as derived equivalences in complete generality. In other words, the action of $\mathrm{SL}_2(\mathbb{Z})$ on the derived category of a (Jacobian) elliptic surface extends to general noncommutative rational surfaces with $K^2 = 0$. We are also able to understand derived equivalences more generally; if

$K^2 > 0$, or if a certain root system is finite, then we can completely characterize the derived equivalences between the surfaces we have constructed. In addition, the same ideas that let us control derived equivalences are even more powerful when applied to abelian equivalences, allowing us to completely control those.

Although we used the notional association with difference equations to motivate the construction via difference operators and various of the results above, most of the results discussed above have little to do with that interpretation. There are several resulting loose ends in connecting things back to equations, the most important of which being the relation between the various Poisson maps between moduli spaces and "isomonodromy" transformations. From one perspective, this is fairly straightforward: we develop enough machinery below to let the computations of such transformations in [47] to carry over *mutatis mutandis*; in particular, we find that the guesses made there for how to incorporate q in the actions of those transformations are indeed correct descriptions of the corresponding actions from noncommutative geometry. This is still essentially algebraic, apart from one issue, namely what it means for such a transformation to preserve "monodromy" in the elliptic setting.

One answer comes from [50], which shows that given a q-difference equation with coefficients meromorphic on $\mathbb{C}^*/\langle p\rangle$, there is a natural way to associate an equivalence class of p-difference equations on $\mathbb{C}^*/\langle q\rangle$, which one can view as the monodromy of the equation. The noncommutative geometry perspective suggests an improvement: we are able to show that there is a version of this construction that gives an equivalence between the corresponding categories of coherent sheaves disjoint from C. (Curiously, unlike the differential case, there is no need to introduce perverse sheaves to make this an abelian equivalence.) In addition, this "elliptic Riemann-Hilbert functor" behaves well in holomorphic families, and thus induces biholomorphic maps between the relevant moduli spaces. The construction of this functor is based on a general construction of a sheaf of solutions associated to any q-difference equation with meromorphic coefficients, which may be of independent interest.

Moreover, the elliptic Riemann-Hilbert functor commutes with the various Cremona transformations, while twisting by a line bundle has no effect on the image of the functor (so preserves "monodromy" as required). The commutation with Cremona transformations is mostly straightforward, but there is one tricky case. The fact that $\mathbb{P}^1 \times \mathbb{P}^1$ is a Hirzebruch surface in two different ways means that a sheaf on $\mathbb{P}^1 \times \mathbb{P}^1$ has *two* interpretations as a difference equation. In terms of difference operators, this corresponds to a sort of formal Fourier transform, swapping certain multiplication operators and difference operators. To show that this operation respects the elliptic Riemann-Hilbert functor, we need to remove the word "formal", and instantiate the operation as an actual integral transform. This leads to a couple of interesting features: one is that we obtain an action of a certain large Coxeter group via integral transformations, and can recover several known elliptic hypergeometric identities as relations between different reduced words. The other is that the geometry tells us that any rigid sheaf corresponds to a -2-curve, which is thus related by a sequence of Cremona transformations to the -2-curve on the Hirzebruch surface F_2. Reversing these operations and using the analytic interpretation, we find that we can construct a solution to any rigid equation via a sequence of integral transformations. The case with only one integration step corresponds to the higher-order elliptic beta integral [60, 51], showing that these indeed satisfy rigid equations.

The plan of the paper is as follows. In Section 2, using results of [47] as motivation, we construct the $\mathbb{Z}^2$-algebras corresponding to our noncommutative deformations of $\mathbb{P}^1 \times \mathbb{P}^1$; Section 3 then shows that this is a flat deformation (and constructs important symmetries), and Section 4 shows that there is still a well-behaved category of coherent sheaves. Section 4 also considers the "missing" operators: in order to have a flat deformation, we had to omit some of the morphisms between line

bundles, but it turns out that the omitted morphisms still correspond to difference operators.

Section 5 gives the analogous construction for the Hirzebruch surface F_1. The presence of a -1-curve on this surface leads to two important features: one is that the orthogonal $\mathbb{Z}$-subalgebra is a noncommutative deformation of $\mathbb{P}^2$, while the other is that there is a *natural* first-order difference equation associated to every object in the category. If we gauge by solutions of those equations, we obtain an equivalent category in which now every operator that appears has elliptic function (as opposed to meromorphic theta function) coefficients. Moreover, this allows us to describe the generators in a purely algebraic way, which is key to making our later constructions work for general genus 1 curves.

Section 6 constructs surfaces with $K^2 = 7$. The idea here is that such a surface is a blowup of both kinds of Hirzebruch surface, and thus to understand such surfaces one must merely understand elementary transformations. This leads to a description of the Hom spaces in the $\mathbb{Z}^3$-algebra as intersections of Hom spaces in the F_0 and F_1 algebras, mediated by a suitable gauge transformation; we can also give a description via generators that lets us extend the algebraic construction.

Section 7 finally extends the construction to general $K^2 < 7$. Here the main difficulty is that we have been unable to give an explicit description of the relevant Hom spaces: there are difficulties when one blows up the same point multiple times. What saves the construction is that we are hoping for a *flat* deformation, which it turns out allows us to reconstruct the deformation from the generic fiber, where the Hom spaces are easy to describe. Moreover, the generic fiber has all the symmetries we are hoping for, making it easy to extend the symmetries to the full family, and letting us reduce the proof of flatness to the fundamental chamber for the appropriate Coxeter group, where it is reasonably straightforward.

Section 8 is a mainly technical section regarding when the multiplication maps in the $\mathbb{Z}^{m+2}$-algebra are surjective, in preparation for showing that the various notions of "torsion" module agree. This is shown in Section 9, which in particular shows that our construction agrees with the construction à la Van den Bergh, as well as satisfying symmetries under the appropriate Coxeter group of Cremona transformations. Section 9 also gives an algorithm for computing Hom spaces in the saturated $\mathbb{Z}^{m+2}$-algebra (i.e., maps between line bundles on the noncommutative scheme). Section 10 then shows that the resulting noncommutative schemes satisfy the axioms of [14]. Finally, Sections 11-13 discuss moduli spaces, derived equivalences, and analytic issues, as described above.

There are also two appendices: Although we mostly avoid working with presentations in favor of simply specifying generators inside a large ambient algebra of difference operators, there are some low-order cases in which there are particularly nice generators and relations. Indeed, for both $\mathbb{P}^1 \times \mathbb{P}^1$ and two-point blowups thereof, not only can we give an explicit presentation for generic parameters, but that presentation actually takes the form of a Gröbner basis. (This is in sharp contrast to the elliptic noncommutative $\mathbb{P}^2$, for instance.) Appendix A considers the general question of when relations of the given form for $\mathbb{P}^1 \times \mathbb{P}^1$ satisfy the Gröbner basis property, and shows that this leads to a nice integrable system on the $\mathbb{P}^{15}$ parametrizing the relations. Appendix B considers the $K^2 = 6$ case, and shows that it extends to a similar deformation of the blowup of $\mathbb{P}^n$ in the $n+1$ coordinate vectors. This leads to a noncommutative deformation of $(\mathbb{P}^1)^n$ for general n, and a number of results essentially stating that noncommutative surfaces of the kind we consider can be embedded as subschemes of such blowups.

Further directions

We intend to extend most of the above results in future work. If one allows the anticanonical curve to degenerate, then the strong flatness result of Section 7 already fails to hold for commutative surfaces. (More precisely, this happens whenever the anticanonical curve has a smooth rational component of self-intersection < -2, see [47].) As a result, the present approach would become considerably more complicated, but the resulting insights are still applicable, and indeed we will be

able to extend nearly all of the results of Sections 9–12 to the general case, as well as (to a lesser extent) to more general noncommutative surfaces [44].

This is not to say that there is no generalization of the flatness result. Many of the various explicit difference operators appearing below have natural multivariate analogues, and it turns out that this extends to give a multivariate analogue of the family of $\mathbb{Z}^{m+2}$-algebras constructed below, satisfying the same Coxeter group symmetry and a slightly weaker version of flatness. These considerations will also lead to an alternate construction of the present $\mathbb{Z}^{m+2}$-algebras as "spherical algebras" of an "elliptic double affine Hecke algebra" constructed in [46].

Another class of open problems involves morphisms between our noncommutative surfaces. Although we have fully determined the *isomorphisms* between such surfaces, we have little understanding of higher-degree morphisms. For instance, if one of our surfaces is equipped with a faithful action of a finite group, then presumably the quotient by that action will not only exist, but be another noncommutative rational surface. Moreover, any morphism of noncommutative rational surfaces would give rise to a corresponding pair of functors acting on difference equations, some of which will presumably have natural interpretations. We can also consider sheaves which are fixed by the action of an automorphism; examples of this include difference equations in which the shift matrix is in the orthogonal or symplectic group, leading to the question of what the noncommutative framework might tell us about cocycles in general semisimple groups over the field of elliptic functions.

There are also a number of open questions suggested by the results of Appendix B. Indeed, the blowups of $\mathbb{P}^n$ constructed there presumably give rise to smooth proper noncommutative n-folds in the sense of [14], and similarly for the $(\mathbb{P}^1)^n$ constructed there. Moreover, the latter can presumably be viewed as a $\mathbb{P}^1$-bundle over $(\mathbb{P}^1)^{n-1}$, suggesting that the construction of [65] can be extended to $\mathbb{P}^1$-bundles over noncommutative schemes. The embeddings of noncommutative Hirzebruch surfaces in noncommutative $\mathbb{P}^n$s also raises the question of whether other very ample divisor classes on noncommutative rational surfaces give rise to such embeddings (for different deformations of $\mathbb{P}^n$). A possibly related question is whether there is an analogue of Castelnuovo-Mumford regularity for general very ample divisor classes.

Notation

Although for most of the paper, we will be working with general (algebraic) genus 1 curves, it will be convenient towards the beginning and necessary in Section 13 to work with analytic curves. It is not only traditional in the theory of elliptic special functions but convenient when developing the theory of sheaves of solutions to work with such curves in *multiplicative* notation. That is, rather than view a typical analytic elliptic curve as $\mathbb{C}/\Lambda$ for some lattice Λ, we instead exponentiate away one of the periods to obtain an expression as $\mathbb{C}^*/\langle p\rangle$ for $|p| < 1$.[2] This makes modular transformations relatively inaccessible, but those are of only minor interest below. One perhaps somewhat odd effect this has on the notation below is that we will continue to use multiplicative notation for the group law of algebraic elliptic curves, and will similarly express divisors in multiplicative notation. Thus if $x \in C$, then $[x]^l$ denotes the divisor with order l at x and 0 elsewhere, and similarly for products and ratios. (This is not as unnatural as it may appear to those coming from an algebraic perspective, as it means that if f and g are two elliptic functions, then $\operatorname{div}(fg) = \operatorname{div}(f)\operatorname{div}(g)$.)

As usual in the analytic setting, to specify an elliptic function, it is simplest to express it as a ratio of products of quasiperiodic functions. The primary such function in the multiplicative

[2]Again, since q is traditionally used to denote the shift in the difference equation, it is traditional in work on elliptic special functions to use p to denote the nome of the elliptic curve.

notation is the function

$$\theta_p(z) := \prod_{0\le i}(1-p^i z)(1-p^{i+1}/z), \tag{1.5}$$

defined for $|p|<1$, which satisfies the symmetries $\theta_p(1/z) = -z^{-1}\theta_p(z)$ and $\theta_p(p/z) = \theta_p(z)$ as well as the quasiperiodicity

$$\theta_p(pz) = -z^{-1}\theta_p(z). \tag{1.6}$$

More generally, a "p-theta function" is a holomorphic function f satisfying a quasiperiodicity relation of the form $f(pz) = Az^{-k}f(z)$. A particularly important special case is that of symmetric theta functions: a "$BC_1(\eta)$-symmetric p-theta function of degree d" is a holomorphic function f such that $f(\eta/z) = f(z)$ and $f(pz) = (\eta/pz^2)f(z)$. (By standard convention a theta function is always holomorphic unless specifically stated otherwise; thus a meromorphic function satisfying the above conditions would be called a "meromorphic ($BC_1(\eta)$-symmetric) theta function".)

Another important function is the "elliptic Gamma function" [52]:

$$\Gamma_{p,q}(z) := \prod_{0\le j,k} \frac{1-p^{j+1}q^{k+1}/z}{1-p^j q^k z}, \tag{1.7}$$

defined for $|p|$,$|q|<1$, which satisfies the difference equations

$$\Gamma_{p,q}(qz) = \theta_p(z)\Gamma_{p,q}(z),$$
$$\Gamma_{p,q}(pz) = \theta_q(z)\Gamma_{p,q}(z),$$

as well as the reflection principle $\Gamma_{p,q}(pq/z) = \Gamma_{p,q}(z)$.

We also define

$$\theta_p(z;q)_k := \prod_{0\le j<k} \theta_p(q^j z), \tag{1.8}$$

defined for $|p|<1$, $q\in\mathbb{C}^*$, $k\in\mathbb{Z}$, where by convention we interpret such products for negative k by taking

$$\prod_{0\le j<k} f_j = \prod_{k\le j<0} f_j^{-1}. \tag{1.9}$$

Note that if $|q|<1$, then

$$\theta_p(z;q)_k = \frac{\Gamma_{p,q}(q^k z)}{\Gamma_{p,q}(z)}. \tag{1.10}$$

We will mainly be using it in this form; i.e., the extension to general $q\in\mathbb{C}^*$ of the cocycle corresponding to $\Gamma_{p,q}$.

Finally, by convention, we take the natural numbers $\mathbb{N}$ to include 0.

Acknowledgements. The author would like to thank A. Borodin, P. Etingof, A. Okounkov, and M. Van den Bergh for helpful conversations. This work was partially supported by grants from the National Science Foundation, DMS-1001645 and DMS-1500806.

2 Difference equations as modules

An (analytic) *elliptic difference equation* is a formal equation of the form

$$v(qz) = A(z)v(z) \tag{2.1}$$

where $A(z)$ is a matrix of meromorphic theta functions, i.e., $A(pz) = \alpha z^k A(z)$ for some $\alpha \in \mathbb{C}^*$, $k \in \mathbb{Z}$, and $|q| < 1$. While general equations of this type can be related to sheaves on noncommutative surfaces (specifically noncommutative $\mathbb{P}^1$-bundles over elliptic curves [44]), for present purposes, we are only interested in a special case. Indeed, the elliptic difference equations arising in the theory of elliptic special functions are invariably *symmetric*, in the sense that $A(1/qz) = A(z)^{-1}$ (note that this forces $k = 0$). The reason for this is quite simple: for the most part, the solutions we have in mind for those equations satisfy an additional condition, namely that $v(1/z) = v(z)$. The symmetry condition is merely the obvious compatibility condition, extending the difference equation to a 1-cocycle for the infinite dihedral group. (Indeed, it follows from results of [42] that the symmetry condition is precisely that needed to obtain a fundamental matrix of symmetric meromorphic solutions to the difference equation, see also the discussion in Section 13.2 below.)

As noted in [47], a symmetric elliptic difference equation admits a factorization: by the matrix version of Hilbert's Theorem 90, A can always be written in the form

$$A(z) = B(1/qz)^{-t} B(z)^t, \tag{2.2}$$

where $B(pz) \propto B(z)$. This factorization is nonunique for two reasons. The first is that if $B(pz) = \beta z^l B(z)$, then $A(pz) = \beta^2 (pq)^{-l} A(z)$, so that there are multiple choices for which line bundle the coefficients of B are sections of. Of course, when the coefficients of A are actually elliptic, this is no issue, and in any case, it turns out that there is usually an obvious choice of multiplier for B. The more serious issue is that we can right multiply B by any meromorphic theta function matrix C such that $C(1/qz) = C(z)$; this, as noted in [47] can be fixed by replacing B by a holomorphic map of vector bundles, and insisting that its domain be maximal. To be precise, B should be a morphism

$$B : \pi_1^* V \to \pi_0^* W, \tag{2.3}$$

where V and W are vector bundles on $\mathbb{P}^1$, and $\pi_i : \mathbb{C}/\langle p \rangle \to \mathbb{P}^1$ is the quotient by $z \mapsto 1/q^i z$. (Note that the matrices C we may multiply by are themselves pulled back through π_1.) Since vector bundles on $\mathbb{P}^1$ are sums of line bundles, we can replace this by the requirement that B have meromorphic coefficients with

$$B_{ij}(pz) = \beta z^l (1/pz^2)^{d_{1i}} (1/pqz^2)^{-d_{2j}} B_{ij}(z). \tag{2.4}$$

In the canonical factorization of A, $d_{1i} \equiv 0$, and $\sum_j d_{2j}$ is as large as possible. The result is then unique up to automorphisms of the domain bundle V.

We thus arrive at equations of the form

$$B(z)^t v(z) - B(1/qz)^t v(qz) = 0 \tag{2.5}$$

and want to interpret these as modules over a suitable algebra. Of course, as we have already noted, we cannot expect this to work, so instead should obtain modules over a suitable $\mathbb{Z}^2$-algebra. The coefficients of the above matrix equation are sums of operators of the form

$$v_i(z) \mapsto B_{ij}(z) v_i(z) - B_{ij}(1/qz) v_i(qz) \tag{2.6}$$

Of course, these operators are only really determined up to scalar multiplication. We can be guided here by the fact that since we want to consider symmetric meromorphic solutions, our operators need to act nicely on the space of symmetric meromorphic functions. If $v_i(z) = v_i(1/z)$, then

$$B_{ij}(z) v_i(z) - B_{ij}(1/qz) v_i(qz) = B_{ij}(z) v_i(z) - B_{ij}(1/qz) v_i(1/qz) \tag{2.7}$$

is antisymmetric under $z \mapsto 1/qz$. We should thus at the very least multiply by an antisymmetric function so that the operator preserves symmetry. When v_i is holomorphic near a point, the output of the above operator will be holomorphic, and will vanish at fixed points of $z \mapsto 1/qz$. This suggests instead considering the operator

$$v_i(z) \mapsto \frac{B_{ij}(z)v_i(z) - B_{ij}(1/qz)v_i(qz)}{z^{-1}\theta_p(qz^2)}. \tag{2.8}$$

This operator takes meromorphic functions invariant under $z \mapsto 1/z$ to meromorphic functions invariant under $z \mapsto 1/qz$. At this point, we could if desired choose a square root of q, and use it to shift the output to make it invariant under $z \mapsto 1/z$. This is somewhat clumsy, and in particular introduces the question of whether the result would depend on the choice of square root. Since we are looking for a category in any event, we may as well simply allow the symmetry involution to vary with the object in the category. Once we do this, there is nothing particularly special about the involution $z \mapsto 1/z$, so we may as well replace it by $z \mapsto q\eta/z$. (The appearance of q here is to simplify later formulas.)

We thus end up with operators of the form

$$\mathcal{D}_{\eta;q;p}(b) := \frac{z}{\theta_p(z^2/\eta)}(b(z) - b(\eta/z)T) \tag{2.9}$$

with $b(z)$ a (holomorphic) theta function, and T the operator $Tf(z) = f(qz)$. Note that left-multiplication by symmetric theta functions changes the multiplier of $b(z)$ by a power of η/pz^2. These operators take functions invariant under $z \mapsto q\eta/z$ to functions invariant under $z \mapsto \eta/z$. For the composition of such operators to make sense, we need two things. First, the symmetry conditions must match: we must only consider compositions $\mathcal{D}_{\eta/q,q}(b_1)\mathcal{D}_{\eta,q}(b_2)$. Second, in such a composition, the coefficient of T is a sum of two terms, and we will naturally want both terms to be sections of the same line bundle. If $b_1(pz) = \beta_1 z^{l_1} b_1(z)$, $b_2(pz) = \beta_2 z^{l_2} b_2(z)$, then this condition forces

$$\frac{\beta_1^2(\eta/pq)^{l_1}}{q^4\beta_2^2(q\eta/p)^{l_2}} = 1. \tag{2.10}$$

This expresses β_1/β_2 as a square root; the choice of square root is suggested by the fact that the coefficient of T is invariant under $z \mapsto \eta/qz$, and both terms have a denominator factor $z^{-1}\theta_p(qz^2/\eta)$ which we would prefer to cancel. This factor introduces four poles in the fundamental annulus; two of those are automatically cancelled, but the poles at $\pm\sqrt{\eta/pq}$ give conditions

$$\frac{\beta_1(\pm\sqrt{\eta/pq})^{l_1-l_2}}{q^{2+l_2}\beta_2} = 1. \tag{2.11}$$

In particular, we are also forced here to have $l_1 - l_2$ even, lest the two conditions force opposite signs. This will of course split our construction into two cases, corresponding to even and odd Hirzebruch surfaces.

In general, we will want a $\mathbb{Z}^2$-algebra in which the objects are linear combinations of symbols s and f, in such a way that f records the extent to which we have multiplied by symmetric theta functions, and s records the degree of the difference operator. We will in particular want a category in which $\mathrm{Hom}(ds + d_1'f, (d+1)s + d_2'f)$ consists of operators $\mathcal{D}_{\eta;q;p}(b)$ with b a theta function with suitable multiplier; if we insist that $\mathrm{Hom}(ds + d_1'f, ds + d_2'f)$ consist of multiplication by suitably symmetric theta functions, then this together with the above consistency condition will determine all of the multipliers once we have chosen the multiplier corresponding to $\mathrm{Hom}(0, s)$. Moreover, there will always be some k such that the line bundle corresponding to $\mathrm{Hom}(0, s + kf)$ has degree

1 or 2, and we can perform a change of basis in the group of objects so that $k = 0$. The degree 1 case will correspond to odd Hirzebruch surfaces, while the degree 2 case will correspond to even Hirzebruch surfaces. We will focus on the latter case for the moment, as it has a very important additional symmetry.

In the above considerations, we needed to assume $|q| < 1$, or more precisely that q is not a root of unity, so that the operators act faithfully on meromorphic functions. However, we may ignore that representation, and simply view T as a formal symbol satisfying the *operator* equation $Tf(z) = f(qz)T$, where we view $f(z)$ as the operator of multiplication by a meromorphic function. More generally, we obtain an algebra of meromorphic q-difference operators, which are polynomials

$$\sum_{0\leq k\leq d} c_k(z)T^k \tag{2.12}$$

with each $c_k(z)$ a meromorphic function, and with multiplication given by

$$(\sum_{0\leq k\leq d} c_k(z)T^k)(\sum_{0\leq k\leq d'} c'_k(z)T^k) = \sum_{0\leq l\leq d+d'} \sum_{\max(0,l-d')\leq k\leq \min(l,d)} c_k(z)c'_{l-k}(q^k z)T^l. \tag{2.13}$$

We may now define a $\mathbb{C}$-linear category $\mathcal{S}_{\eta,\eta';q;p}$ as follows. First, define a $\mathbb{Z}s + \mathbb{Z}f$-algebra $\mathcal{M}er\mathcal{D}iff_q$ in which each Hom space is the algebra of meromorphic q-difference operators, with the obvious composition. Then $\mathcal{S}_{\eta,\eta';q;p}$ is the smallest subcategory of $\mathcal{M}er\mathcal{D}iff_q$ having the same objects, and containing the following morphisms:

- If $g(z)$ is a $BC_1(q^{1-d}\eta)$-symmetric theta function of degree 1, then

$$g(z) \in \mathcal{S}_{\eta,\eta';q;p}(ds + d'_f, ds + (d'+1)f). \tag{2.14}$$

- If $h(z)$ is a $BC_1(q^{1-d'}\eta')$-symmetric theta function of degree 1, then

$$\mathcal{D}_{q^{-d}\eta;q;p}(h) \in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + d'f). \tag{2.15}$$

- If $b(z)$ is holomorphic such that $b(pz) = (q^{-d-d'+1}\eta\eta'/p^2z^4)b(z)$, then

$$\mathcal{D}_{q^{-d}\eta;q;p}(b) \in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + (d'+1)f). \tag{2.16}$$

We will show below that this family of categories is flat: the dimension of any given Hom space is independent of the parameters. Moreover, when $q = 1$ and $\eta/\eta' \notin p^{\mathbb{Z}}$, the category is isomorphic to the subcategory of $\mathrm{Coh}_{\mathbb{P}^1\times\mathbb{P}^1}$ with objects $\mathcal{O}(ds + d'f)$ (where $\mathcal{O}(s)$ is the class of a section and $\mathcal{O}(f)$ the class of a fiber relative to one of the two rulings of $\mathbb{P}^1 \times \mathbb{P}^1$), and the natural "swap rulings" automorphism of $\mathbb{P}^1 \times \mathbb{P}^1$ extends to the full family of categories.

This construction is essentially invariant under translation of the group of objects, a fact we may use to largely restrict our attention to morphisms from the object 0.

Lemma 2.1. *As spaces of operators,*

$$\mathcal{S}_{\eta,\eta';q;p}(d_1s + d'_1f, d_2s + d'_2f) = \mathcal{S}_{q^{-d}\eta,q^{-d'}\eta';q;p}((d_1-d)s + (d'_1-d')f, (d_2-d)s + (d'_2-d')f), \tag{2.17}$$

and this gives an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{q^{-d}\eta,q^{-d'}\eta';q;p} \tag{2.18}$$

which on objects takes $d_1s + d'_1f$ to $(d_1 - d)s + (d'_1 - d')f$.

Proof. Indeed, it suffices to prove this for the spaces of generators, where it is immediate. □

One might worry that we are missing some first-order operators in the above construction, since we only included those in which $b(z)$ comes from a line bundle of degree 2 or 4. In fact, we are essentially missing no such operators. For convenience, we only consider morphisms starting from 0, as the remaining cases will follow by translation invariance.

Lemma 2.2. *If $b(z)$ is a theta function with*

$$b(pz) = (q\eta'/pz^2)(\eta/pz^2)^{d'}b(z), \tag{2.19}$$

$d' \geq 0$, then

$$\mathcal{D}_{\eta;q;p}(b) \in \mathcal{S}_{\eta,\eta';q;p}(0, s + d'f). \tag{2.20}$$

Proof. If $d' \in \{0, 1\}$, there is nothing to prove, as the given operators are already generators. We proceed by induction, and assume we have already proved the result for $d' - 1$. Now, if $g(z)$ is a $BC_1(\eta)$-symmetric theta function of degree 1, then

$$g\mathcal{D}_{\eta;q;p}(b') = \mathcal{D}_{\eta;q;p}(gb'); \tag{2.21}$$

it will thus suffice to show that the space of theta functions with multiplier

$$b(pz)/b(z) = (q\eta'/pz^2)(\eta/pz^2)^{d'} \tag{2.22}$$

is spanned by those functions of the form $g(z)b'(z)$ where g, b' are theta functions with multipliers

$$g(pz)/g(z) = \eta/pz^2, \qquad b'(pz)/b'(z) = (q\eta'/pz^2)(\eta/pz^2)^{d'-1}. \tag{2.23}$$

This is a special case of the following lemma, which we state in an algebraic way (and will use extensively). □

Lemma 2.3. *Let C be a smooth genus 1 curve over a field k, and let $\mathcal{L}$, $\mathcal{L}'$ be line bundles on C with $\deg(\mathcal{L}), \deg(\mathcal{L}') \geq 2$ and either $\mathcal{L} \ncong \mathcal{L}'$ or $\deg(\mathcal{L}) > 2$. Then the multiplication map*

$$\Gamma(C; \mathcal{L}) \otimes \Gamma(C; \mathcal{L}') \to \Gamma(C; \mathcal{L} \otimes \mathcal{L}') \tag{2.24}$$

is surjective.

Proof. Since $\deg(\mathcal{L}) \geq 2$, we may choose two global sections f_1, $f_2 \in \Gamma(\mathcal{L})$ that have no zeros in common. As a result, they define a short exact sequence

$$0 \to \mathcal{L}' \otimes \mathcal{L}^{-1} \xrightarrow{(f_1, f_2)} \mathcal{L}' \oplus \mathcal{L}' \xrightarrow{(f_2, -f_1)} \mathcal{L} \otimes \mathcal{L}' \to 0 \tag{2.25}$$

If $\mathcal{L}' \otimes \mathcal{L}^{-1}$ is acyclic, then the second morphism remains surjective on taking global sections, and the claim follows. Since the symmetric argument applies whenever $\mathcal{L} \otimes \mathcal{L}'^{-1}$ is acyclic, we have proved the desired result whenever $\mathcal{L} \ncong \mathcal{L}'$.

If $\mathcal{L} \cong \mathcal{L}'$, our hypotheses ensure that $\deg(\mathcal{L}) > 2$. But then for any point $x \in C$ (the claim is geometric, so we may extend scalars as necessary to ensure a point exists), the Lemma gives a surjection

$$\Gamma(\mathcal{L}(-x)) \otimes \Gamma(\mathcal{L}) \to \Gamma(\mathcal{L}^{\otimes 2}(-x)). \tag{2.26}$$

It follows that for any point $x \in C$, the image of

$$\Gamma(\mathcal{L}) \otimes \Gamma(\mathcal{L}) \to \Gamma(\mathcal{L}^{\otimes 2}) \tag{2.27}$$

contains every section vanishing at x. Since any section vanishes at *some* point, the claim follows. □

Remark. For an analytic version of essentially the same argument, see [43].

The only missing operators are those with $d' < 0$; since the multiplier of a theta function is determined (up to a power of p) by its zeros, we find that such operators can only exist if $d' = -1$ and $q\eta' \in p^{\mathbb{Z}}\eta$. While we could include such operators, the result would necessarily fail to be a flat deformation; moreover, even including such operators would fail to include all the operators that "ought" to be there, see Section 4 below.

To complete the connection between difference equations and the category $\mathcal{S}_{\eta,\eta';q;p}$, it remains only to construct a module for each difference equation, in such a way that we can recover the difference equations (and, ideally, the *solutions* of the difference equation) from the module. To this end, we introduce a natural representation of $\mathcal{S}_{\eta,\eta';q;p}$: each morphism of $\mathcal{S}_{\eta,\eta';q;p}$ is a difference operator with meromorphic coefficients, and thus acts on the space of all meromorphic functions. More precisely, we should impose the appropriate symmetry conditions, and thus obtain a module $\mathcal{M}er$ over $\mathcal{S}_{\eta,\eta';q;p}$ (i.e., a functor from $\mathcal{S}_{\eta,\eta';q;p}$ to $\mathbb{C} - \mathcal{V}ect$) as follows: $\mathcal{M}er(ds + d's)$ is the space of meromorphic functions $f(z)$ satisfying $f(\eta/q^{d-1}z) = f(z)$, while if

$$\mathcal{D} = \sum_{0\le k\le d} c_k(z)T^k \in \mathcal{S}_{\eta,\eta';q;p}(d_0 s + d_0' f, (d_0 + d)s + (d_0' + d')f), \tag{2.28}$$

then

$$\mathcal{D}\cdot f(z) = \sum_{0\le k\le d} c_k(z) f(q^k z). \tag{2.29}$$

For any $d_0, d_0' \in \mathbb{Z}$, we also define a module $P_{d_0 s + d_0' f}$ by

$$P_{d_0 s + d_0' f}(ds + d'f) := \mathcal{S}_{\eta,\eta';q;p}(-d_0 s - d_0' f, ds + d'f), \tag{2.30}$$

with the obvious multiplication. Note that the usual Yoneda construction gives

$$\mathrm{Hom}(P_{d_0 s + d_0' f}, P_{ds + d'f}) \supset \mathcal{S}(-ds - d'f, -d_0 s - d_0' f). \tag{2.31}$$

Theorem 2.4. *Given any symmetric elliptic difference equation*

$$v(qz) = A(z)v(z) \tag{2.32}$$

with $A(pz) = (q\eta'/\eta)^2 A(z)$, there is a corresponding module M_A over $\mathcal{S}_{\eta,\eta';q;p}$ such that the space of homomorphisms $\mathrm{Hom}(M_A, \mathcal{M}er)$ *can be naturally identified with the space of meromorphic vectors such that*

$$v(qz) = A(z)v(z), \qquad v(1/z) = v(z). \tag{2.33}$$

Proof. Factor the equation as

$$B(\eta/z)^t v(qz) = B(z)^t v(z) \tag{2.34}$$

as above, and interpret each operator

$$\mathcal{D}_{\eta;q;p}(B_{ij}(z)) \tag{2.35}$$

as an element of

$$\mathcal{S}_{\eta,\eta';q}(0, s + d'f) \subset \mathrm{Hom}(P_{-s-d'f}, P_0). \tag{2.36}$$

(This fails if $q\eta' = \eta$ and some coefficients of B are constant; to fix this, we may multiply such operators by a basis of the space of $BC_1(\eta)$-symmetric theta functions. We omit the details, since this problem goes away once we have introduced the saturated category.)

This allows us to interpret B as a morphism of the form

$$\bigoplus_{1\le k\le n} P_{-s+d'_k f} \to P_0^n. \tag{2.37}$$

Let M_A be the cokernel of this morphism. The morphisms from P_0 to $\mathcal{M}er$ can be naturally identified with meromorphic functions invariant under $z \mapsto q\eta/z$, and thus morphisms from M_A to $\mathcal{M}er$ can be identified with vectors of such functions. The vectors that arise are those such that the corresponding morphism from P_0^n becomes trivial on the domain of B; in other words, such that

$$B(\eta/z)^t v(qz) = B(z)^t v(z). \tag{2.38}$$

□

Remark. Since B is essentially unique once we maximize $\sum_k d'_k$ (automorphisms of the vector bundle clearly lift to automorphisms of the above domain), M_A as constructed is itself unique up to isomorphism. A similar construction applies to a symmetric equation in which

$$A_{ij}(pz) = (q\eta'/\eta)^2 (q\eta/pz^2)^{e_i - e_j} A_{ij}(z), \tag{2.39}$$

the only modification being to replace P_0^n by a sum $\oplus_{1\le i\le n} P_{e_i f}$.

In general, translation invariance tells us that we are missing operators whenever $\eta/\eta' \in p^{\mathbb{Z}} q^{\mathbb{Z}}$. (We will see later that this is the only instance in which we can enlarge the Hom spaces in a natural way.) On the other hand, most of the time, our spaces of generators are redundant. Indeed, the same argument tells us that the natural composition map

$$\mathcal{S}(s, s+f) \otimes \mathcal{S}(0, s) \to \mathcal{S}(0, s+f) \tag{2.40}$$

is surjective unless $\eta/q\eta' \in p^{\mathbb{Z}}$, when we only obtain operators $\mathcal{D}_{\eta;q;p}(b)$ where $z^l b$ is $BC_1(\eta)$-symmetric of degree 2 for some integer l (depending on the power of p in the ratio of multipliers). Similarly,

$$\mathcal{S}(f, s+f) \otimes \mathcal{S}(0, f) \to \mathcal{S}(0, s+f) \tag{2.41}$$

is surjective unless $q\eta/\eta' \in p^{\mathbb{Z}}$, when the leading coefficient is essentially a $BC_1(q\eta)$-symmetric theta function. In particular, at least one of the two compositions is surjective unless $q^2 \in p^{\mathbb{Z}}$ and $q\eta/\eta' \in p^{\mathbb{Z}}$. If neither map is surjective but $q \notin p^{\mathbb{Z}}$, we obtain the span of the spaces of $BC_1(\eta)$-symmetric and $BC_1(q\eta)$-symmetric theta functions of degree 2; these are distinct codimension 1 subspaces of the relevant space of all theta functions, so span. In other words, we find that $\mathcal{S}_{\eta,\eta';q;p}$ is generated in degrees f, s unless $q, \eta/\eta' \in p^{\mathbb{Z}}$. (This corresponds to the commutative Hirzebruch surface F_2).

3 Flatness (even Hirzebruch case)

As mentioned above, one reason for the above definition of $\mathcal{S}_{\eta,\eta';q;p}$ is that the resulting family of $\mathbb{Z}^2$-algebras is flat. To be precise, we have the following.

Theorem 3.1. *For all $\eta, \eta', q, p \in \mathbb{C}^*$, $|p| < 1$, and any integers d, d',*

$$\dim(\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f)) = \max(d+1, 0)\max(d'+1, 0). \tag{3.1}$$

Remark. This is, of course, trivially true for $d < 0$ or $d' < 0$, since the only generators are in the positive quadrant.

To show this, we will prove both a lower and an upper bound on the dimensions of the Hom spaces. The lower bound, of course, simply involves constructing sufficiently many linearly independent difference operators in the space. The key idea will be to understand the image and kernel of the leading coefficient map $\mathcal{D} \mapsto [T^0]\mathcal{D}$. Note that this is not just linear: it clearly takes composition in $\mathcal{S}$ to multiplication of (meromorphic) theta functions.

Lemma 3.2. *If $d' \geq 0$, then*

$$\dim([T^0]\mathcal{S}_{\eta,\eta';q;p}(0,d'f)) = d'+1; \tag{3.2}$$

if $d \geq 0$, then

$$\dim([T^0]\mathcal{S}_{\eta,\eta';q;p}(0,ds)) = d+1. \tag{3.3}$$

Finally, if $d, d' > 0$, then

$$\dim([T^0]\mathcal{S}_{\eta,\eta';q;p}(0,ds+d'f)) = 2d+2d'. \tag{3.4}$$

Proof. A section of $\mathcal{S}_{\eta,\eta';q;p}(0,d'f)$ is a linear combination of products of generators of degree f, and as such is a $BC_1(q\eta)$-symmetric theta function of degree d'. Any such function factors into degree 1 functions, and thus we obtain the full space, of dimension $d'+1$. A similar calculation applies to the $d'=0$ case.

For $d, d' > 0$, the leading coefficient is a section of a bundle of degree $2d+2d'$, so we need merely show that all sections arise. This follows from Lemma 2.3 by induction: if it holds for $ds+d'f$, then it holds for $(d+1)s+d'f$ and $ds+(d'+1)f$ upon multiplication by generators of degrees s, f respectively. We thus reduce to the case $d=d'=1$, where it follows by inspection of the generators of that degree. □

This already gives us a tight lower bound if d or d' is 0 or 1 (so that we reduce to showing that in that case, the operator is determined by its leading coefficient). The first case where we obtain an additional operator is $d=d'=2$, where it turns out that the operator T can be obtained. This was already shown in [43], but we will give a different argument for a slightly stronger fact.

Lemma 3.3. *For any η, η', q, p, there exists an expansion*

$$T = \sum_i \mathcal{D}_{\eta/q;q;p}(b_{i1})\mathcal{D}_{\eta;q;p}(b_{i2}), \tag{3.5}$$

where for each i, b_{i1}, b_{i2} are theta functions with

$$b_{i1}(pz) = (\eta\eta'/p^2qz^4)b_{i1}(z),$$
$$b_{i2}(pz) = (q\eta\eta'/p^2z^4)b_{i2}(z).$$

Moreover, for any $v \in \mathbb{C}^*$, there exists such an expansion with $b_{i1}(v)=0$ for all i.

Proof. Let $b_i(x,y) := b_{i1}(x)b_{i2}(y)$, and note that we can write

$$\begin{aligned}\mathcal{D}_{\eta/q;q;p}(b_{i1})\mathcal{D}_{\eta;q;p}(b_{i2}) = {} & \frac{z^2 b_i(z,z)}{\theta_p(z^2/\eta, qz^2/\eta)} + \frac{qz^2 b_i(\eta/qz, \eta/qz)}{\theta_p(z^2/\eta, q^2z^2/\eta)}T^2 \\ & - \frac{z^2 b_i(z,\eta/z)}{\theta_p(z^2/\eta, qz^2/\eta)}T - \frac{qz^2 b_i(\eta/qz, qz)}{\theta_p(qz^2/\eta, q^2z^2/\eta)}T.\end{aligned} \tag{3.6}$$

This is, of course, linear in b_i, so a similar formula holds for the sum of these product operators, in terms of the function $b(x,y) = \sum_i b_i(x,y)$. This function can be any element of the tensor product of the two spaces of theta functions; i.e., the only constraint on b is that it be holomorphic and have multipliers

$$b(px,y) = (\eta\eta'/p^2qx^4)b(x,y),$$
$$b(x,py) = (q\eta\eta'/p^2y^4)b(x,y).$$

Note that the constraint that $b_{i1}(v) = 0$ for all i translates to the condition that $b(v,y) = 0$.

Now, consider the function

$$b(x,y) = \frac{\theta_p(y/x, \eta\eta'/qxyvv', x/v, x/v', y/w, q^2vv'/wy)}{\theta_p(\eta/qvv', \eta'/qvv', qv/w, qv'/w)} \tag{3.7}$$

for v', $w \in \mathbb{C}^*$ such that the denominator is nonzero. This certainly satisfies $b(v,y) = 0$, so the corresponding operator has an expansion as required. The coefficients of T^0 and T^2 are 0, since $b(z,z) = 0$, while the coefficient of T^1 simplifies via the addition law. We find that the resulting operator is $q^{-1}\eta T$, so the result follows. □

Remark. Note that as long as v' and w are such that the denominator does not vanish, we obtain an expansion satisfying $b_{i1}(v) = b_{i1}(v') = b_{i2}(w) = 0$ for all i.

In particular, we find that

$$\dim(\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f)) \geq \dim([T^0]\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f)) + \dim(T\mathcal{S}_{\eta,\eta';q;p}(0, (d-2)s + (d'-2)f)), \tag{3.8}$$

which gives the desired lower bound by induction.

For the upper bound, we note that since our category is given as a generated subcategory of a flat family, we have semicontinuity: specializing the parameters can only make the Hom spaces smaller. We thus need only prove the upper bound for generic parameters. Thus, suppose $\eta/\eta' \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$. As we have already noted, this means that the generators of degree $s + f$ are redundant, so we have precisely four generators from each object. To be precise, for each $d \in \mathbb{Z}$, let g_{1d}, g_{2d} be a basis of the space of $BC_1(q^{1-d}\eta)$-symmetric theta functions of degree 1, and for each $d' \in \mathbb{Z}$, let $h_{1d'}$, $h_{2d'}$ be a basis of the space of $BC_1(q^{1-d'}\eta)$-symmetric theta functions of degree 1. Then the category $\mathcal{S}_{\eta,\eta';q;p}$ is generated by the elements

$$x_i(d,d') := g_{id} \in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f) \tag{3.9}$$

and

$$y_i(d,d') := \mathcal{D}_{q^{-d}\eta;q;p}(h_{id'}) \in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + d'f). \tag{3.10}$$

To prove the desired upper bound, we will need to consider the relations that these generators satisfy. Luckily, it will turn out that quadratic relations suffice.

There are three cases to consider for quadratic relations: relations of degree $2f$, $s + f$, and $2s$. The first set is quite simple: since elements of degree f are multiplication operators, they simply commute:

$$x_1(d,d'+1)x_2(d,d') = x_2(d,d'+1)x_1(d,d'). \tag{3.11}$$

The degree $s + f$ is only slightly more complicated: since

$$x_i(d+1,d')y_j(d,d') = \mathcal{D}_{q^{-d}\eta;q;p}(g_{i(d+1)}h_{jd'})$$
$$y_i(d,d'+1)x_j(d,d') = \mathcal{D}_{q^{-d}\eta;q;p}(h_{i(d'+1)}g_{jd}),$$

we find that

$$\sum_{ij} \alpha_{ij} y_i(d, d'+1) x_j(d, d') = \sum_{ij} \beta_{ij} x_i(d+1, d') y_j(d, d') \tag{3.12}$$

iff

$$\sum_{ij} \alpha_{ij} h_{i(d'+1)} g_{jd} = \sum_{ij} \beta_{ij} g_{i(d+1)} h_{jd'}. \tag{3.13}$$

The constraint $\eta/\eta' \notin p^{\mathbb{Z}} q^{\mathbb{Z}}$ implies that the four functions on either side are linearly independent, and thus we obtain four relations of the form

$$y_i(d, d'+1) x_i(d, d') = \sum_{kl} M_{ijkl}(d, d') x_i(d+1, d') y_j(d, d'). \tag{3.14}$$

The relations of degree $2s$ require somewhat more thought, but it turns out that they are in fact quite nice.

Lemma 3.4. *If h_1, h_2 are $BC_1(\eta')$-symmetric theta functions of degree 1, then*

$$\mathcal{D}_{\eta/q;q;p}(h_1)\mathcal{D}_{\eta;q;p}(h_2) = \mathcal{D}_{\eta/q;q;p}(h_2)\mathcal{D}_{\eta;q;p}(h_1). \tag{3.15}$$

Proof. As above, we can express the difference of the two operators in terms of the function

$$h(x, y) = h_1(x)h_2(y) - h_2(x)h_1(y) \tag{3.16}$$

By symmetry considerations, we find

$$h(x, y) \propto x^{-1}\theta_p(x/y, xy/\eta'), \tag{3.17}$$

from which it is easy to see that the operator vanishes. □

In other words, there is one degree $2s$ relation, to the effect that

$$y_1(d+1, d')y_2(d, d') = y_2(d+1, d')y_1(d, d'). \tag{3.18}$$

Now, given any homogeneous linear combination of products of generators, we may use the degree $s+f$ relations to move the "x" generators to the right, then use the degree $2f$ and $2s$ relations to sort the "x" and "y" generators amongst themselves. It follows that $\mathcal{S}_{\eta,\eta';q;p}(0, ds+d'f)$ is spanned by products of the form

$$y_1^k y_2^{d'-k} x_1^l x_2^{d'-l}, \tag{3.19}$$

where by y_i^m we mean the appropriate product of the form

$$y_i^m := y_i(d_0+m-1, d_0')y_i(d_0+m-2, d_0') \cdots y_i(d_0, d_0') \tag{3.20}$$

The upper bound follows, and concludes the proof of Theorem 3.1. □

Since the lower and upper bounds agree, the proof of the lower bound gives us the following useful fact.

Corollary 3.5. *If $\mathcal{D} \in \mathcal{S}_{\eta,\eta';q;p}(0, ds+d'f)$ is such that $[T^0]\mathcal{D} = 0$, then*

$$T^{-1}\mathcal{D} \in \mathcal{S}_{\eta,\eta';q;p}(0, (d-2)s + (d'-2)f), \tag{3.21}$$

$$\mathcal{D}T^{-1} \in \mathcal{S}_{\eta,\eta';q;p}(2s+2f, ds+d'f). \tag{3.22}$$

Similarly, the fact that the bounds agree implies that the relations considered in the proof of the upper bound actually give a presentation. More than that, the relations actually give a quadratic Gröbner basis (with respect to a suitable monomial ordering), since we can put elements into canonical form with moves consisting only of replacing subwords with linear combinations of smaller subwords. This presumably implies that when $\eta/\eta' \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$, the category $\mathcal{S}_{\eta,\eta';q;p}$ is Koszul in some suitable sense. Similarly, apart from a presumably mild nondegeneracy condition, this implies that $\mathcal{S}_{\eta,\eta';q;p}$ satisfies a form of Artin-Schelter regularity. We will show in Appendix A below that this essentially characterizes the category: any $\mathbb{Z}^2$-algebra with a similar presentation is either isomorphic to $\mathcal{S}_{\eta,\eta';q;p}$ or to a degeneration thereof.

An even more striking consequence of this presentation stems from the fact that the degree f generators and the degree s generators satisfy very similar relations. Indeed, each generator is identified with a suitable theta function, and the relations are just the quadratic relations satisfied by those functions. (Compare the abstract description of the three-generator Sklyanin algebra [3].) In particular, we find that when $\eta/\eta' \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$, we obtain an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{\eta',\eta;q;p}, \tag{3.23}$$

which on objects takes $ds + d'f$ to $d's + df$. It turns out that the constraint on η/η' can be removed here.

Theorem 3.6. *There is an isomorphism*

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{\eta',\eta;q;p}, \tag{3.24}$$

which on objects takes $ds + d'f$ to $d's + df$, and on generators takes

$$\begin{aligned} g(z) &\in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f) \\ \mathcal{D}_{q^{-d}\eta;q;p}(h) &\in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + d'f) \\ \mathcal{D}_{q^{-d}\eta;q;p}(b) &\in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + (d'+1)f) \end{aligned}$$

to

$$\begin{aligned} \mathcal{D}_{q^{-d'}\eta';q;p}(g) &\in \mathcal{S}_{\eta',\eta;q;p}(d's + df, (d'+1)s + df), \\ h &\in \mathcal{S}_{\eta',\eta;q;p}(d's + df, d's + (d+1)f), \\ \mathcal{D}_{q^{-d'}\eta';q;p}(b) &\in \mathcal{S}_{\eta',\eta;q;p}(d's + df, (d'+1)s + (d+1)f) \end{aligned}$$

respectively.

Proof. This transformation is clearly an involution, so we need only show that it extends to a homomorphism. In other words, given any relation satisfied by the generators of $\mathcal{S}_{\eta,\eta';q;p}$, we need to show that their images in $\mathcal{S}_{\eta',\eta;q;p}$ satisfy the same relation. The key point is that for any given relation, this is a *closed* condition, and thus we need only prove it for $\eta/\eta' \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$. For relations involving only the degree f and s generators, this follows by checking the above quadratic relations, which as we have noted simply reduces to ordinary multiplication.

It thus remains only to show that the homomorphism defined by the given action on degree f and s generators acts in the stated way on degree $s + f$ generators. By our genericity assumption, there is a basis of such generators consisting of elements of the form

$$g\mathcal{D}_{q^{-d}\eta;q;p}(h) = \mathcal{D}_{q^{-d}\eta;q;p}(gh) \in \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + (d'+1)f) \tag{3.25}$$

But this transforms to

$$\mathcal{D}_{q^{-d'}\eta'}(g)h = \mathcal{D}_{q^{-d'}\eta'}(gh) \tag{3.26}$$

as required. ∎

Since this involution exchanges multiplication and difference operators, we call it the "Fourier transform". (In fact, under a suitable limit, it degenerates to the usual Fourier-Laplace transform, see [44].) We should also note that the Fourier transform essentially preserves the elements T; to be precise, if we express T in terms of degree $s+f$ generators as above, it is then straightforward to show that the transform takes

$$q^{-d}\eta T \in \mathcal{S}_{\eta,\eta';q;p}(ds+d'f,(d+2)s+(d'+2)f) \tag{3.27}$$

to $q^{-d'}\eta' T$.

A similar argument shows that any family of isomorphisms between leading coefficient categories $[T^0]\mathcal{S}_{\eta,\eta';q;p}$ extends to isomorphisms between the full categories. Indeed, for $\eta/\eta' \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$, we find that both $\mathcal{S}_{\eta,\eta';q;p}$ and $[T^0]\mathcal{S}_{\eta,\eta';q;p}$ have the same generators, and the same low-degree relations; moreover, the low-degree relations suffice to give a presentation of $\mathcal{S}_{\eta,\eta';q;p}$. Any isomorphism of leading coefficient categories takes generators to generators preserving low-degree relations, so extends. We may then use flatness to relax the constraint on the parameter: any family of isomorphisms extends to the closure of the family in parameter space. The morphisms in $[T^0]\mathcal{S}_{\eta,\eta';q;p}$ can be identified with (holomorphic) theta functions of multiplier

$$\frac{z^{-2d_2-2d_2'}q^{d_2+d_2'-d_2d_2'}(\eta/p)^{d_2'}(\eta'/p)^{d_2}}{z^{-2d_1-2d_1'}q^{d_1+d_1'-d_1d_1'}(\eta/p)^{d_1'}(\eta'/p)^{d_1}}, \tag{3.28}$$

so any transformation of theta functions that gives the correct multiplier and respects multiplication will produce an isomorphism of $\mathcal{S}$.

For instance, if $g(z)$ is a theta function with the above multiplier, then $g(\alpha z)$ has essentially the same multiplier, except with η,η' replaced by $\alpha^{-2}\eta,\alpha^{-2}\eta'$. We thus conclude that there is an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{\alpha^{-2}\eta,\alpha^{-2}\eta';q;p} \tag{3.29}$$

for any $\alpha \in \mathbb{C}^*$, acting as the identity on objects. (Note that when $\alpha=-1$, this isomorphism becomes an automorphism.) The action of this isomorphism on difference operators is simple: just conjugate by the operator T_α that translates by α. We similarly obtain an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{\beta^2/\eta,\beta^2/\eta';1/q;p}, \tag{3.30}$$

conjugating by the reflection $g(z)\mapsto g(\beta/z)$. There is also the trivial gauge transformation that acts on leading coefficients by

$$g(z)\mapsto (C(d_2s+d_2'f)/C(d_1s+d_1'f))g(z), \tag{3.31}$$

where $C:\mathbb{Z}^2\to\mathbb{C}^*$ is any map.

A slightly more subtle transformation multiplies $g(z)$ by $z^{d_2-d_1}$; this produces an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{\eta,p\eta';q;p}. \tag{3.32}$$

This acts on difference operators by taking

$$\sum_{0\le k\le d_2-d_1} c_k(z)T^k \in \mathcal{S}_{\eta,\eta';q;p}(d_1s+d_1'f,d_2s+d_2'f) \tag{3.33}$$

to

$$\sum_{0\le k\le d_2-d_1} z^{d_2-d_1-2k}\eta^k q^{-k(k+d_1-1)}c_k(z)T^k; \tag{3.34}$$

it is easy to see that this acts as an automorphism on $\mathcal{M}er\mathcal{D}iff_q$, and takes generators to generators. We can also view this as a gauge transformation, but with nonconstant functions: for any $a \in \mathbb{C}^*$, if we define functions

$$F_{ds+d'f}(z) = q^{d(d-1)/2}(-\eta a)^{-d}\theta_q(az, aq^{1-d}\eta/z), \tag{3.35}$$

then

$$\sum_{0\leq k\leq d_2-d_1} z^{d_2-d_1-2k}\eta^k q^{-k(k+d_1-1)}c_k(z)T^k = F_{d_2s+d_2'f}(z)^{-1}\sum_{0\leq k\leq d_2-d_1} c_k(z)T^k F_{d_1s+d_1'f}(z). \tag{3.36}$$

If we conjugate the above isomorphism by the Fourier transform, we obtain an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{p\eta,\eta';q;p}; \tag{3.37}$$

note that together with the previous isomorphism, this lets us produce automorphisms from translations by $\pm\sqrt{p}$, in addition to the translation by -1 automorphism considered above. (It follows from the elliptic construction below that this can also be described as a nonconstant gauge transformation, we omit the details.) There is also an isomorphism which acts on leading coefficients by

$$g(z) \mapsto z^{-d_2d_2'+d_1d_1'}g(z), \tag{3.38}$$

and gives an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}_{\eta,\eta';pq;p}. \tag{3.39}$$

This invariance under p-shifting the parameters is, of course, not unexpected, and simply reflects (as does modular invariance, the details of which we omit) the underlying algebraic nature of the construction. This will be clearer in the case of the noncommutative F_1 considered below.

There is one more isomorphism of interest coming from the above construction. The morphisms in $[T^0]\mathcal{S}$ are morphisms between line bundles on $\mathbb{C}/\langle p\rangle$, and as such we can obtain a *contravariant* isomorphism by simply dualizing those bundles. In this way, we obtain an isomorphism

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}^{\mathrm{op}}_{\eta,\eta';1/q;p}, \tag{3.40}$$

which acts on objects by $ds + d'f \mapsto (2-d)s + (2-d')f$, and on leading coefficients as the identity. (This action on objects is motivated by the fact that $2s + 2f$ is the anticanonical divisor.) This involution has a particularly nice interpretation at the level of difference operators. There is a natural contravariant involution between $\mathcal{M}er\mathcal{D}iff_q$ and $\mathcal{M}er\mathcal{D}iff_{1/q}$, acting via

$$\sum_{0\leq k\leq d} c_k(z)T_q^k \mapsto \sum_{0\leq k\leq d} T_{1/q}^k c_k(z) = \sum_{0\leq k\leq d} c_k(q^{-k}z)T_{1/q}^k. \tag{3.41}$$

We can think of this as a sort of formal adjoint: if $\mathcal{D}$ is a meromorphic q-difference operator, and f, g are any meromorphic functions, then we have the formal integral identity

$$\int (\mathcal{D}f)(z)g(z)\mu(dz) = \int f(z)(\mathcal{D}^{\mathrm{ad}}g)(z)\mu(dz) \tag{3.42}$$

for any measure μ invariant under q-shifts, with $\mathcal{D}^{\mathrm{ad}}$ given by the above involution. (Of course, this only applies as an actual identity of integrals when the various terms on both sides are integrable, a relatively unlikely event, and only determines $\mathcal{D}^{\mathrm{ad}}$ when q is non-torsion; for our purposes, though, there is no problem with doing purely formal manipulations of integrals.) While this naïve formal

adjoint does not preserve the form of morphisms in $\mathcal{S}$, we can fix things by a suitable gauge transformation. We thus replace the naïve formal adjoint on $\mathcal{M}er\mathcal{D}iff$ by the adjoint formally satisfying

$$\int z^{-1}\theta_p(q^{d_2-1}z^2/\eta)g(z)(\mathcal{D}f)(z)\mu(dz) = \int z^{-1}\theta_p(q^{d_1-1}z^2/\eta)(\mathcal{D}^{\mathrm{ad}}g)(z)f(z)\mu(dz), \tag{3.43}$$

for $\mathcal{D} \in \mathcal{M}er\mathcal{D}iff_q(d_1 s + d_1' f, d_2 s + d_2' f)$, in which case

$$\mathcal{D}^{\mathrm{ad}} \in \mathcal{M}er\mathcal{D}iff_{1/q}((2-d_2)s + (2-d_2')f, (2-d_1)s + (2-d_1')f). \tag{3.44}$$

Explicitly we have

$$(\sum_{0\leq k} c_k(z)T_q^k)^{\mathrm{ad}} = \sum_{0\leq k} \frac{q^k\theta_p(q^{d_2-2k-1}z^2/\eta)}{\theta_p(q^{d_1-1}z^2/\eta)}c_k(q^{-k}z)T_{1/q}^k. \tag{3.45}$$

This is still contravariant, and acts correctly on generators, so restricts to an involution

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}^{\mathrm{op}}_{\eta,\eta';1/q;p} \tag{3.46}$$

as required.

The notion of formal adjoint also gives an alternate interpretation of the Fourier transform. We would like the Fourier transform to arise as conjugation by a suitable family of operators $F_{\eta,\eta';q;p}$, so that

$$\mathcal{D}F_{q^{-d_1}\eta,q^{-d_1'}\eta';q;p} = F_{q^{-d_2}\eta,q^{-d_2'}\eta';q;p}\hat{\mathcal{D}}. \tag{3.47}$$

If we take the ansatz that $F_{\eta,\eta';q;p}$ should be an integral operator

$$(F_{\eta,\eta'}g)(z) = \int K_{\eta,\eta'}(z,w)g(w)\mu(dw) \tag{3.48}$$

(which, again, will have definitional issues as anything other than a formal object), then (3.47) becomes an equation for $K_{\eta,\eta'}(z,w)$, comparing the action of $\mathcal{D}$ in the z variables to the action of the naïve formal adjoint of $\hat{\mathcal{D}}$ in the w variables. Since we know the Fourier transform is an isomorphism, we need only consider the equations coming from generators. The equation for $\mathcal{D} \in \mathcal{S}'(0,s)$ is given by

$$\mathcal{D}_{\eta;q;p}(\theta(z/a, q\eta'/az))_z K_{\eta,\eta';q;p}(z,w) = \theta(w/a, q\eta'/aw)K_{\eta/q,\eta';q;p}(z,w). \tag{3.49}$$

Setting $a = w$ gives a first-order equation in $K_{\eta,\eta';q;p}(z,w)$ as a function of z, implying

$$K_{\eta,\eta';q;p}(z,w) = \theta_q(z, q\eta/z)\Gamma_{p,q}(z/w, q\eta/zw, zw/q\eta', w\eta/\eta' z)G_{\eta,\eta'}(z,w), \tag{3.50}$$

where $G_{\eta,\eta'}(z,w)$ is q-elliptic in z. Taking $a = z$ then gives the first-order equation

$$G_{\eta/q,\eta'}(z,w) = \theta_p(\eta/q\eta')G_{\eta,\eta'}(z,w), \tag{3.51}$$

allowing us to eliminate the dependence in η (again up to an elliptic factor). A similar calculation using the other generators lets us solve for the w and η'-dependence, giving us

$$G_{\eta,\eta'}(z,qw) = -\frac{\theta_p(w^2/q\eta')}{\theta_p(\eta'/qw^2)}G_{\eta,\eta'}(z,w), \tag{3.52}$$

leading us to the following expression for the kernel (up to an indeterminable factor which is q-elliptic in each of η, η', z, w):

$$K_{\eta,\eta';q;p}(z,w) = \frac{\theta_q(z,q\eta/z)}{\theta_q(w,q\eta'/w)} \frac{\Gamma_{p,q}(z/w, q\eta/zw, zw/q\eta', w\eta/\eta' z)}{\Gamma_{p,q}(w^2/q\eta', q\eta'/w^2, \eta/\eta')} \tag{3.53}$$

Apart from the q-theta function factors, this is essentially the same as the kernel for an integral operator considered in [62].

The fact that we need a formal interpretation above is analogous to the fact that the corresponding identity for the Laplace transform (swapping differentiation and multiplication) only holds up to boundary terms from the requisite integration by parts. Here, if we replace the formal invariant measure with an actual measure, we will have error terms corresponding to the failure to be invariant under q-shifts. If one can control those terms (say if the integral is a contour integral, and the region between the contour and its q shift contains no poles), then the formal calculation gives rise to actual difference equations.

As an example, consider the operator

$$\mathcal{D} = \mathcal{D}_{\eta;q;p}(z^{-1}\theta_p(a_1 z, a_2 z, a_3 z, q a_1 a_2 a_3 \eta\eta'/z)) \in \mathcal{S}_{\eta,\eta';q;p}(0, s+f). \tag{3.54}$$

The equation $\hat{\mathcal{D}}g(w) = 0$ is straightforward to solve:

$$g(w) = \theta_q(w, q\eta'/w) \frac{\prod_{1\le i\le 3} \Gamma_{p,q}(a_i w, q a_i \eta'/w)}{\Gamma_{p,q}(q^2 a_1 a_2 a_3 \eta\eta'/w, q a_1 a_2 a_3 w\eta/b_1)}. \tag{3.55}$$

Of course, this is only defined up to an overall q-elliptic factor, but with this particular choice, we find (generically) that for a suitable choice of contour C, the formal equation

$$\mathcal{D} \int_C K_{\eta,\eta'}(z,w) g(w) \frac{dw}{w} = \int_C K_{\eta,\eta'}(z,w) (\hat{\mathcal{D}} g)(w) \frac{dw}{w} = 0 \tag{3.56}$$

is actually valid analytically. The corresponding evaluation of the integral arising from this first-order equation (up to an overall elliptic factor) is essentially just the elliptic beta integral of [60], and the proof is essentially the same.

A similar calculation shows that the above integral operator takes solutions of equations of degree $s + df$ to solutions of equations of degree $ds + f$, recovering the fact that order $d-1$ elliptic beta integrals satisfy symmetric elliptic difference equations of degree d [51]. One can also make the case $2s + 2f$ precise, see [53]. More generally, one expects that if $\mathcal{D}$ and $\hat{\mathcal{D}}$ are Fourier transforms, then it should be possible to span the space of functions annihilated by $\mathcal{D}$ by integrals of the form

$$\int_C K_{\eta,\eta'}(z,w) g(w) \frac{dw}{w} \tag{3.57}$$

with $\hat{\mathcal{D}}g = 0$. We will prove a precise version of this (subject to some technical assumptions) below.

4 Torsion modules and sheaves I

One disadvantage of dealing with a $\mathbb{Z}^d$-algebra is that, as with multigraded commutative algebras, the usual dictionary between modules and sheaves breaks down. An obvious solution is to choose a coset of some subgroup $\mathbb{Z} \subset \mathbb{Z}^d$ and restrict the category to the objects in that coset to obtain a $\mathbb{Z}$-algebra. This certainly gives a notion of sheaf, but the resulting category of sheaves can depend

significantly on the choice of subgroup (e.g., in the commutative setting, the Proj of the graded algebra associated to a divisor depends in particular on whether the divisor is ample), and it can be tricky to determine when two such choices give rise to equivalent categories of sheaves.

In the case of the $\mathbb{Z}^2$-algebra $\mathcal{S}_{\eta,\eta';q;p}$, we may take an alternate approach. First, note that in the commutative degeneration, the divisor class $ds + d'f$ is *always* ample if $d' > d > 0$; if $d' \leq d$, it fails to be ample on F_2 (it contains or is orthogonal to the -2-curve, which has class $s - f$), while if $d = 0$, the divisor class maps the surface to $\mathbb{P}^1$. We will show that any of these always ample classes gives rise to the same category of sheaves on $\mathcal{S}_{\eta,\eta';q;p}$.

The key ingredient in the construction of the category of sheaves on a $\mathbb{Z}$-algebra is the notion of a torsion module, one in which each element is annihilated by all morphisms of sufficiently high degree. For $\mathcal{S}_{\eta,\eta';q;p}$, we replace that notion by the following.

Definition 4.1. Let $\mathcal{M}$ be a module over $\mathcal{S}_{\eta,\eta';q;p}$. A homogeneous element $v \in \mathcal{M}(d_0 s + d_0' f)$ is torsion (of order $d_1 s + d_1' f$) if there exists some $d_1' \geq d_1 \geq 0$ such that for all $d' \geq d \geq 0$, we have

$$\mathcal{S}_{\eta,\eta';q;p}(d_0 s + d_0' f, (d_0 + d_1 + d)s + (d_0' + d_1' + d')f)v = 0. \tag{4.1}$$

The module $\mathcal{M}$ is torsion iff all homogeneous elements of $\mathcal{M}$ are torsion.

Note that a finite linear combination of torsion elements is torsion: simply maximize $d_1' - d_1$, d_1 over the orders of the elements in the linear combination.

We will then define a sheaf on $\mathcal{S}_{\eta,\eta';q;p}$ to be an object of the quotient of $\mathcal{S}_{\eta,\eta';q;p}$ by the subcategory of torsion modules. For this to make sense, we need to know that this is a Serre subcategory.

Lemma 4.1. *Let $\mathcal{M}$ be a module over $\mathcal{S}_{\eta,\eta';q;p}$. An element $v \in \mathcal{M}(0)$ is torsion of order $ds + d'f$ with $d' \geq d \geq 0$ iff*

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f)v = 0. \tag{4.2}$$

Proof. A simple inductive argument reduces to showing that if

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f)v = 0, \tag{4.3}$$

then

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + (d'+1)f)v = \mathcal{S}_{\eta,\eta';q;p}(0, (d+1)s + (d'+1)f)v = 0, \tag{4.4}$$

and thus it will suffice to show that

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + (d'+1)f) = \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) \tag{4.5}$$

and

$$\mathcal{S}_{\eta,\eta';q;p}(0, (d+1)s + (d'+1)f) = \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f). \tag{4.6}$$

The construction of the lower bound on the Hom space dimensions shows in particular that whenever $d' \geq d \geq 0$,

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) = \mathcal{S}_{\eta,\eta';q;p}(ds + df, ds + d'f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + df), \tag{4.7}$$

and thus

$$\begin{aligned}
&\mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) \\
&\quad = \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(ds + df, ds + d'f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + df) \\
&\quad = \mathcal{S}_{\eta,\eta';q;p}(ds + df, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + df) \\
&\quad = \mathcal{S}_{\eta,\eta';q;p}(0, ds + (d'+1)f)
\end{aligned} \tag{4.8}$$

as required. Similarly, the adjoint gives

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) = \mathcal{S}_{\eta,\eta';q;p}((d' - d)f, ds + d'f)\mathcal{S}_{\eta,\eta';q;p}(0, (d' - d)f), \tag{4.9}$$

from which the second claim follows. □

Corollary 4.2. *If $d'_1 \geq d_1 \geq 0$ and $d'_2 \geq d_2 \geq 0$, then*

$$\mathcal{S}_{\eta,\eta';q;p}(0, (d_1 + d_2)s + (d'_1 + d'_2)f) = \mathcal{S}_{\eta,\eta';q;p}(d_1 s + d'_1 f, (d_1 + d_2)s + (d'_1 + d'_2)f)\mathcal{S}_{\eta,\eta';q;p}(0, d_1 s + d'_1 f). \tag{4.10}$$

Corollary 4.3. *Let d, d' be integers with $d' > d > 0$. Then an element $v \in M(0)$ is torsion iff for some $k > 0$,*

$$\mathcal{S}_{\eta,\eta';q;p}(0, k(ds + d'f))v = 0. \tag{4.11}$$

Proof. The "if" direction has already been shown; for the only if direction, note that if v is torsion of order $d_1 s + d'_1 f$, then there exists $k > 0$ such that $k(d' - d) \geq d'_1 - d_1$ and $d \geq kd_1$, and thus $kd' - d'_1 \geq kd - d_1 \geq 0$, so that v is annihilated by morphisms of degree $k(ds + d'f)$ as required. □

Proposition 4.4. *The subcategory of torsion modules is a Serre subcategory of the category of all $\mathcal{S}_{\eta,\eta';q;p}$-modules.*

Proof. Since "torsion" is defined in terms of elements, it is preserved as a property of modules by taking submodules. In addition, the image of a torsion element in a quotient module is clearly torsion, so any quotient of a torsion module is torsion. Finally, suppose we have a short exact sequence

$$0 \to T_1 \to M \to T_2 \to 0 \tag{4.12}$$

with T_1, T_2 torsion. For any element $v \in M(0)$, say, its image in T_2 is torsion, and thus there exists $d'_2 \geq d_2 \geq 0$ such that

$$\mathcal{S}_{\eta,\eta';q;p}(0, d_2 s + d'_2 f)v \subset T_1. \tag{4.13}$$

Since this Hom space is finite, we can uniformly choose $d'_1 \geq d_1 \geq 0$ such that

$$\mathcal{S}_{\eta,\eta';q;p}(d_2 s + d'_2 f, (d_1 + d_2)s + (d'_1 + d'_2)f)w = 0 \tag{4.14}$$

for all

$$w \in \mathcal{S}_{\eta,\eta';q;p}(0, d_2 s + d'_2 f)v. \tag{4.15}$$

But then

$$\begin{aligned} &\mathcal{S}_{\eta,\eta';q;p}(0, (d_1 + d_2)s + (d'_1 + d'_2)f)v \\ &\quad = \mathcal{S}_{\eta,\eta';q;p}(d_2 s + d'_2 f, (d_1 + d_2)s + (d'_1 + d'_2)f)\mathcal{S}_{\eta,\eta';q;p}(0, d_2 s + d'_2 f)v \\ &\quad = 0, \end{aligned} \tag{4.16}$$

so that v is torsion as required. □

Remark. This also follows from the second corollary. Indeed, the second corollary tells us that we obtain the same notion of "torsion" upon restricting to any subgroup $\mathbb{Z}(ds + d'f)$ with $d' > d > 0$, and thus we may reduce to known facts about $\mathbb{Z}$-algebras (which by the first corollary are generated in degree 1).

Of course for $\mathbb{P}^1 \times \mathbb{P}^1$, there are far more ample bundles than just those with $d' > d > 0$, and this remains true for the noncommutative deformation.

Proposition 4.5. *Fix integers $d, d' > 0$. If $\eta' \notin p^{\mathbb{Z}} q^{\mathbb{Z}} \eta$, then an element $v \in M(0)$ is torsion iff for some $k > 0$,*

$$\mathcal{S}_{\eta,\eta';q;p}(0, k(ds + d'f))v = 0. \tag{4.17}$$

Proof. We first note that if $d, d' \geq 0$, then

$$\mathcal{S}_{\eta,\eta';q;p}(0, (d+1)s + d'f) = \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (d+1)s + d'f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f), \tag{4.18}$$

and

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + (d'+1)f) = \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f). \tag{4.19}$$

By the Fourier transform, it suffices to prove the latter. If $d' \geq d$, we have already shown this; if $d' < d$, then again the Fourier transform tells us that

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) = \mathcal{S}_{\eta,\eta';q;p}((d-1)s + d'f, ds + d'f)\mathcal{S}_{\eta,\eta';q;p}(0, (d-1)s + d'f). \tag{4.20}$$

But then

$$\begin{aligned}
&\mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) \\
&= \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}((d-1)s + d'f, ds + d'f)\mathcal{S}_{\eta,\eta';q;p}(0, (d-1)s + d'f) \\
&= \mathcal{S}_{\eta,\eta';q;p}((d-1)s + d'f, ds + (d'+1)f)\mathcal{S}_{\eta,\eta';q;p}(0, (d-1)s + d'f) \\
&= \mathcal{S}_{\eta,\eta';q;p}(0, ds + (d'+1)f),
\end{aligned} \tag{4.21}$$

where the second step uses the fact that $\eta' \notin p^{\mathbb{Z}} q^{\mathbb{Z}} \eta$.

In particular, we find that in the above considerations, we may replace the cone $d' \geq d \geq 0$ with the quadrant $d', d \geq 0$ and obtain the same notions of torsion. The result follows as before. □

In particular, given any symmetric elliptic difference equation, we may turn the corresponding module into a sheaf using the above notions. It turns out that we can still recover the difference equation from this sheaf. This reduces to showing that the module of symmetric meromorphic functions is already saturated with respect to torsion modules: i.e., it contains no torsion elements, and any extension by a torsion module is split. This reduces to the following fact. Let P_0 denote the projective module generated by the object 0; that is, P_0 is the functor

$$P_0(ds + d'f) = \mathcal{S}_{\eta,\eta';q;p}(-ds - d'f, 0), \tag{4.22}$$

with the obvious action.

Proposition 4.6. *Let $I \subset P_0$ be a submodule, and let $I' := T^{-1}(I \cap TP_0)$. Then any morphism $\phi : I \to \mathcal{M}er$ extends uniquely to I'.*

Proof. We certainly have an action of T^{-1} on meromorphic functions, so may define for $\mathcal{D} \in I'$, $\psi(\mathcal{D}) = T^{-1}\phi(T\mathcal{D})$. This is certainly a homomorphism, and if $\mathcal{D} \in I$, then

$$\psi(\mathcal{D}) = T^{-1}\phi(T\mathcal{D}) = T^{-1}T\phi(\mathcal{D}) = \phi(\mathcal{D}), \tag{4.23}$$

so ψ is the desired extension. For uniqueness, note that if $\psi : I' \to \mathcal{M}er$ vanishes on I, then

$$T\psi(\mathcal{D}) = \psi(T\mathcal{D}) = 0 \tag{4.24}$$

for all $\mathcal{D} \in I'$, so that $\psi(\mathcal{D}) = 0$. □

Theorem 4.7. *For any torsion module M,* $\mathrm{Hom}(M, \mathcal{M}er) = \mathrm{Ext}^1(M, \mathcal{M}er) = 0$.

Proof. If $v \in \mathcal{M}er(0)$ is torsion of order $ds + d'f$, then the corresponding morphism $P_0 \to \mathcal{M}er$ restricts to the 0 morphism on the ideal generated by degree $ds + d'f$ elements in P_0. If $d \geq 2$, then

$$T\mathcal{S}_{\eta,\eta';q;p}(0, (d-2)s + (d'-2)f) \subset \mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f), \tag{4.25}$$

and thus the morphism vanishes on the ideal generated by elements of degree $(d-2)s + (d'-2)f$. Thus we may as well assume that $d < 2$; since the ideal contains the ideal generated by elements of degree $(d+1)s + (d'+1)f$, we may further assume that $d = 0$. But since elements of degree $d'f$ act on $\mathcal{M}er$ as multiplication operators, there can certainly be no torsion elements of order $d'f$.

For the extensions, we note that it suffices (using translation invariance) to prove that

$$\mathrm{Ext}^1(P_0/I, \mathcal{M}er) = 0 \tag{4.26}$$

for any submodule I such that P_0/I is torsion. The long exact sequence associated to $0 \to I \to P_0 \to P_0/I \to 0$ gives an exact sequence

$$\mathrm{Hom}(P_0, \mathcal{M}er) \to \mathrm{Hom}(I, \mathcal{M}er) \to \mathrm{Ext}^1(P_0/I, \mathcal{M}er) \to \mathrm{Ext}^1(P_0, \mathcal{M}er) = 0. \tag{4.27}$$

Now, since P_0/I is torsion, I contains the ideal generated by elements of degree $ds + d'f$ for some $d' \geq d \geq 0$, and the same argument allows us to reduce to the case $d = 0$. Again, since then some submodule of I is generated by multiplication operators, the result is immediate. □

In particular, if the module M is a subquotient of the module M' such that the other pieces of the filtration are torsion, then $\mathrm{Hom}(M, \mathcal{M}er) = \mathrm{Hom}(M', \mathcal{M}er)$. Thus passing to sheaves preserves the correspondence with solutions of a difference equation. This is not quite enough to let us reconstruct the equation from the sheaf, but should be analogous to the commutative case, and simply rely on mild properties of the direct and inverse image functors corresponding to the map to $\mathbb{P}^1$. Without this, we need to remember the specific presentation as a quotient of direct sums of the sheaves $\hat{P}_v$ associated to the projective modules P_v (analogues of line bundles).

In particular, we may have $\mathrm{Hom}(M, \mathcal{M}er) = \mathrm{Hom}(M', \mathcal{M}er)$ even if the other pieces of the filtration are non-torsion, as long as they are sufficiently small. The point is that just as we can "divide" meromorphic functions by T, we may also divide them by theta functions. As a result, the extension property for homomorphisms holds for any ideal containing any element of degree a multiple of f. This in particular holds for the ideals $\mathcal{S}_{\eta,\eta',x;q;p}(0, ds + d'f - re_1)$ we will construct below. This implies that extension by an extension of point sheaves has no effect on maps to the sheaf of meromorphic functions; this reflects the fact that such extensions correspond to isomonodromy transformations. (This will be discussed in more detail in Section 13.4 below.)

Now that we have a satisfactory notion of "sheaf", we can address the question of "missing" operators in a more systematic way. To be precise, we wish to extend our $\mathbb{Z}^2$-algebra to a $\mathbb{Z}^2$-algebra $\hat{\mathcal{S}}_{\eta,\eta';q;p}$, in which the morphisms are now given by *sheaf* morphisms. That is

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(d_1 s + d_1' f, d_2 s + d_2' f) = \mathrm{Hom}(\hat{P}_{-d_2 s - d_2' f}, \hat{P}_{-d_1 s - d_1' f}). \tag{4.28}$$

It turns out that the morphisms of $\hat{\mathcal{S}}$ can still be identified with difference operators.

Given two modules M_1, M_2 over $\mathcal{S}_{\eta,\eta';q;p}$, the sheaf morphisms from M_1 to M_2 are limits of module morphisms

$$I_{ds+d'f} M_1 \to M_2 \tag{4.29}$$

where $I_{ds+d'f}$ is the ideal generated by elements of degree $ds + d'f$; the limit is taken with respect to the product partial order on the pair $(d' - d, d)$. Thus when $M_1 = P_v$, we again reduce to a question of when morphisms from such ideals extend.

Lemma 4.8. *Fix integers $d, d' \in \mathbb{Z}$, and a submodule $I \subset P_0$. If $T^k \in I$ for some $k \geq 0$ and $I \not\subset TP_0$, then any morphism $\phi : I \to P_{ds+d'f}$ extends uniquely to P_0.*

Proof. Since $I \not\subset TP_0$, there exists an element $\mathcal{D} \in I$ such that $[T^0]\mathcal{D} \neq 0$. We may then write

$$T^k\phi(\mathcal{D}) = \phi(T^k\mathcal{D}) = (T^k\mathcal{D}T^{-k})\phi(T^k). \tag{4.30}$$

Since the first k coefficients of $T^k\phi(\mathcal{D})$ vanish, and the leading coefficient of $T^k\mathcal{D}T^{-k}$ is nonzero, we conclude that the first k coefficients of $\phi(T^k)$ vanish, and thus that there exists a well-defined element

$$T^{-k}\phi(T^k) \in \mathcal{S}(-ds - d'f, 0) \tag{4.31}$$

But this gives the image of 1 under the desired extension, which is unique since $\mathcal{S}$ is a domain. □

In particular, for a general ideal $I \not\subset TP_0$, to extend a map $\phi : I \to P_v$ to P_0, it will suffice to extend it to $I + TP_0$, since the lemma tells us that the extension from $I + TP_0$ to P_0 exists uniquely. We have a short exact sequence

$$0 \to I \to I + TP_0 \to TP_0/(I \cap TP_0) \to 0, \tag{4.32}$$

and thus we reduce to considering the extension problem from the ideal $I \cap TP_0 \subset TP_0 = P_{2s+2f}$. If we start with the ideal generated by elements of degree $ds + d'f$ with $d' \geq d \geq 2$, this reduces the problem to that for the ideal generated by elements of degree $(d-2)s + (d'-2)f$. As before, this lets us reduce to the case $d = 0$. Of course, we cannot expect this extension problem to be solvable in general, but this tells us that if we saturate with respect to the $d = 0$ extension problem, then the resulting module is saturated.

Given a map $\phi : I \to P_v$ where $I \subset P_0$ is the ideal generated by elements of degree $d'f$, then for any nonzero element g of that degree, we may define a meromorphic difference operator $g^{-1}\phi(g)$. Since

$$g\phi(h) = ghg^{-1}\phi(g) = h\phi(g) = \phi(gh), \tag{4.33}$$

this operator is independent of g. This extends to assign a meromorphic operator to any element of P_0, which on I agrees with ϕ. By the above discussion, this allows us to assign a meromorphic difference operator to any morphism in the saturated category $\hat{\mathcal{S}}$.

Lemma 4.9. *The meromorphic operator $\mathcal{D}$ is in $\hat{\mathcal{S}}_{\eta,\eta';q;p}(-ds - d'f, 0)$ iff for some $k \geq 0$ and every $BC_1(q\eta)$-symmetric theta function g of degree k, $g\mathcal{D} \in \mathcal{S}_{\eta,\eta';q;p}(-ds - d'f, kf)$.*

Lemma 4.10. *If $d \leq 0$ or $d' \geq d - 1$, then*

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds + d'f) = \mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f). \tag{4.34}$$

Proof. If $d < 0$, this is immediate, as no difference operator can have negative order. If $d = 0$, then this follows by observing that the regular representation of the $\mathbb{Z}$-algebra associated to $\mathbb{P}^1$ is already saturated.

Thus suppose $d' \geq d - 1 \geq 0$. Let F be the sheaf on $\mathbb{P}^1$ obtained by restricting P_0 to the $\mathbb{Z}$-algebra with objects $ds + \mathbb{Z}f$. The module is finitely generated, since it is generated by the elements of degree $ds + df$ together with finitely many objects of lower degree in f; as a result, we obtain a coherent sheaf. Moreover, the filtration by powers of T corresponds to a filtration of this coherent sheaf in which the successive quotients are contained in pushforwards of line bundles on $\mathbb{C}^*/\langle p \rangle$. Since we know the Hilbert polynomial of the coherent sheaf, and it agrees with the

sum of the Hilbert polynomials of the line bundles, we find that F is an iterated extension of those bundles.

In particular, we have an extension

$$0 \to TF \to F \to L_0 \to 0 \tag{4.35}$$

where L_0 comes from a line bundle of degree $2d + 2d' > 0$ on our elliptic curve, and is thus acyclic. Thus F is acyclic iff TF is acyclic. In other words, the claim holds for $ds + d'f$ iff it holds for $(d-2)s + (d'-2)f$. The base case is $d \in \{-1, 0\}$, where the claim is immediate. □

Lemma 4.11. *Let $d' \geq 0$ be a nonnegative integer, and let $I \subset P_0$ be the ideal generated by elements of degree $d'f$. Then $T^{\lceil d'/2 \rceil} \in I$ unless $q^l \eta/\eta' \in p^{\mathbb{Z}}$ for some $1 \leq l \leq d'$.*

Proof. It suffices to show that every element of degree $d's + d'f$ is contained in I. We can write any such element as a composition of elements of degree $s + f$, and unless $q\eta/\eta' \in p^{\mathbb{Z}}$, each of those elements of degree $s + f$ can be written as a composition xy with x of degree s and y of degree f. We thus find that

$$\mathcal{S}_{\eta,\eta';q;p}(0, d's+d'f) = \mathcal{S}_{\eta,\eta';q;p}((d'-1)s+d'f, d's+d'f)\mathcal{S}_{\eta,\eta';q;p}(f, (d'-1)s+d'f)\mathcal{S}_{\eta,\eta';q;p}(0, f) \tag{4.36}$$

unless $q\eta/\eta' \in p^{\mathbb{Z}}$. Since

$$\mathcal{S}_{\eta,\eta';q;p}(f, (d'-1)s + d'f) = \mathcal{S}_{\eta,\eta'/q;q;p}(0, (d'-1)s + (d'-1)f), \tag{4.37}$$

the claim follows by induction. □

Corollary 4.12. *Let d, d' be integers with $d > \max(d', 0)$. Then*

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds + d'f) = \mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) \tag{4.38}$$

unless $q^{-l}\eta/\eta' \in p^{\mathbb{Z}}$ for some l with $0 < l < \min(d - d', d + 1)$.

Proof. Any operator

$$\mathcal{D} \in \hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds + d'f) \tag{4.39}$$

becomes a morphism of $\mathcal{S}$ when left-multiplied by any $BC_1(\eta/q^{d-1})$-symmetric theta function of sufficiently large degree. In fact, that degree need be no more than $d - d' - 1$, since past that point the module is already saturated. In other words,

$$\mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d-1)f)\mathcal{D} \in \mathcal{S}_{\eta,\eta';q;p}(0, ds + (d-1)f). \tag{4.40}$$

If

$$\begin{aligned}\mathcal{S}_{\eta,\eta';q;p}(ds + (d-1)f, (2d-1-d')s + (d-1)f)\mathcal{S}_{\eta,\eta';q;p}(ds + d'f, ds + (d-1)f) \\ = \mathcal{S}_{\eta,\eta';q;p}(ds + d'f, (2d-1-d')s + (d-1)f),\end{aligned} \tag{4.41}$$

then $T^{d-d'-1}\mathcal{D} \in \mathcal{S}$, and thus $\mathcal{D} \in \mathcal{S}$ as required.

If $d' < -1$, then $\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) = 0$, and thus if

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds + d'f) \neq \mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) = 0, \tag{4.42}$$

then

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds - f) \neq 0 = \mathcal{S}_{\eta,\eta';q;p}(0, ds - f), \tag{4.43}$$

and thus $q^{-l}\eta/\eta' \in p^{\mathbb{Z}}$ for some $1 \leq l \leq d$. □

Corollary 4.13. *If $\eta \notin p^{\mathbb{Z}} q^{\mathbb{Z}} \eta'$, then $\hat{\mathcal{S}}_{\eta,\eta';q;p} = \mathcal{S}_{\eta,\eta';q;p}$.*

As we will see, the above conditions on η, η' are necessary conditions to have no additional operators in a given saturated Hom space. To see this, we will first need to construct an important case of such a missing operator.

Lemma 4.14. *Fix an integer $d \geq 0$, and let $g_1, \dots, g_d$, $h_1, \dots, h_d$ be nonzero $BC_1(q^{1-d}\eta)$-symmetric theta functions of degree 1. Then*

$$h_d \cdots h_1 \mathcal{D}_{q^{1-d}\eta}(g_1)\mathcal{D}_{q^{2-d}\eta}(g_2)\cdots\mathcal{D}_{\eta}(g_d) = g_d \cdots g_1 \mathcal{D}_{q^{1-d}\eta}(h_1)\mathcal{D}_{q^{2-d}\eta}(h_2)\cdots\mathcal{D}_{\eta}(h_d), \tag{4.44}$$

and thus

$$\mathcal{D}_{d;\eta;q;p} := (g_1 \cdots g_d)^{-1}\mathcal{D}_{q^{1-d}\eta}(g_1)\mathcal{D}_{q^{2-d}\eta}(g_2)\cdots\mathcal{D}_{\eta}(g_d) \in \hat{\mathcal{S}}_{\eta,q^{-d}\eta;q;p}(0, ds - df). \tag{4.45}$$

Proof. Since

$$\mathcal{D}_{q^{k-d}\eta}(g_k)\mathcal{D}_{q^{k+1-d}\eta}(g_{k+1}) = \mathcal{D}_{q^{k-d}\eta}(g_{k+1})\mathcal{D}_{q^{k+1-d}\eta}(g_k), \tag{4.46}$$

we can freely commute the g and h functions amongst themselves. It thus remains to show that we can swap g_1 and h_1. In other words, we need to show

$$h_1 \mathcal{D}_{q^{1-d}\eta}(g_1) = g_1 \mathcal{D}_{q^{1-d}\eta}(h_1). \tag{4.47}$$

Since both sides equal $\mathcal{D}_{q^{1-d}\eta}(g_1 h_1)$, the claim is immediate. □

Remark. Of course, since the category is invariant up to isomorphism under translation of η and η' by powers of p, we obtain such additional elements whenever $q^d\eta' \in p^{\mathbb{Z}}\eta$.

Note in particular that since we have a nontrivial section of $\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, l(s-f))$, we find

$$\begin{aligned}
\dim(\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds + d'f)) &\geq \dim(\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(ls - lf, ds + d'f)) \\
&\geq (d - l + 1)(d' + l + 1) \\
&> (d+1)(d'+1) \\
&= \dim(\mathcal{S}_{\eta,q^{-l}\eta;q;p}(0, ds + d'f))
\end{aligned} \tag{4.48}$$

as long as $0 < l < \min(d - d', d + 1)$. This is already enough to tell us that our condition above was tight.

To do better and determine precisely the structure of $\hat{\mathcal{S}}$ where it differs from $\mathcal{S}$, we need to understand the case $\eta/\eta' = q^l$ better. It turns out that this case has an additional significance.

Lemma 4.15. *Any operator in $\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds + d'f)$ takes $BC_1(q\eta)$-symmetric theta functions of degree $l-1$ to $BC_1(q^{1-d}\eta)$-symmetric theta functions of degree $l + d' - d - 1$.*

Proof. Given a $BC_1(q\eta)$-symmetric theta function f, and an operator $\mathcal{D} \in \hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds + d'f)$, we may compute $\mathcal{D} \cdot f = (\mathcal{D}f) \cdot 1$. In this way, we immediately reduce to the case $l = 1$, and wish to understand the image of 1 under such an operator. The claim is certainly preserved under saturation: if $g\mathcal{D}\cdot 1$ is a $BC_1(q^{1-d}\eta)$-symmetric theta function of degree $d' - d + 1$ for all g symmetric of degree 1, then $\mathcal{D}\cdot 1$ is symmetric of degree $d' - d$. As a result, we may restrict to the case $d' \geq d$.

Now, we have the factorization

$$\hat{\mathcal{S}}_{\eta,\eta/q;q;p}(0, ds + d'f) = \mathcal{S}_{\eta,\eta/q;q;p}(0, ds + d'f) = \mathcal{S}_{\eta,\eta/q;q;p}(ds + df, ds + d'f)\mathcal{S}_{\eta,\eta/q;q;p}(0, ds + df). \tag{4.49}$$

Any operator in $\mathcal{S}_{\eta,\eta/q;q;p}(ds+df, ds+d'f)$ clearly takes constants to $BC_1(q^{1-d}\eta)$-symmetric theta functions of degree $d'-d$, so we may reduce to the case $d'=d$. Since any such operator is a linear combination of products of operators of degree $s+f$, we reduce to that case. That is, we need to know that

$$\frac{(b(z)-b(\eta/z))z}{\theta_p(z^2/\eta)} \tag{4.50}$$

is constant for any holomorphic function b with $b(pz)=(\eta/pz^2)^2 b(z)$. Since this function is elliptic, it suffices to prove it holomorphic, which is easily checked pole-by-pole. □

Since

$$\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds+d'f) = \hat{\mathcal{S}}_{\eta,\eta;q;p}(lf, ds+(d'+l)f), \tag{4.51}$$

and the property $\eta=\eta'$ is invariant under twisting by multiples of $s+f$, we can restate this fact in the following form.

Corollary 4.16. *The $\hat{\mathcal{S}}_{\eta,\eta;q;p}$-module $\mathcal{M}er$ has a natural submodule $\mathcal{S}kl$ such that $\mathcal{S}kl(ds+d'f)$ is the space of $BC_1(q^{1-d}\eta)$-symmetric theta functions of degree $d'-d-1$.*

Remark. If we restrict this module to the $\mathbb{Z}$-algebra $lf+\mathbb{Z}(s+f)$, we may view it as an l-dimensional module over Sklyanin's noncommutative $\mathbb{P}^3$ [58]. This family of representations is precisely the "principal analytic series" constructed in [59]. From a geometric perspective, $\mathcal{S}kl$ is the noncommutative analogue of the sheaf $O_{s-f}(-1)$ on F_2 (that is, the sheaf $O(-1)$ on the section of self-intersection -2).

In particular, if $d'\le d-l$, then any operator in $\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds+d'f)$ *annihilates* all $BC_1(q\eta)$-symmetric theta functions of degree $l-1$, since it is supposed to take such functions to $BC_1(q^{1-d}\eta)$-symmetric theta functions of negative degree. This in particular applies to the operator $\mathcal{D}_{l;\eta;q;p}$. In fact, this is an l-th order operator, and is annihilating an l-dimensional space, so this essentially characterizes $\mathcal{D}_{l;\eta;q;p}$. This is not quite true, as there are some technicalities when q is torsion. We do, however, have the following fact.

Lemma 4.17. *Suppose $l>0$ is a positive integer and that the subgroup of $\mathbb{C}^*/\langle p\rangle$ generated by q has order at least l. If*

$$\mathcal{D} \in \hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds+d'f) \tag{4.52}$$

annihilates every $BC_1(q\eta)$-symmetric theta function of degree $l-1$, then

$$\mathcal{D} \in \hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(l(s-f), ds+d'f)\mathcal{D}_{l;\eta;q;p}. \tag{4.53}$$

Proof. This is clearly preserved by saturation, so we may as well assume $d'\ge d$. If $d>l$, then we find

$$[T^0]\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds+d'f) = [T^0](\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(l(s-f), ds+d'f)\mathcal{D}_{l;\eta;q;p}), \tag{4.54}$$

since both sides surject on to the space of global sections of the same line bundle. We may thus subtract a suitable multiple of $\mathcal{D}_{l;\eta;q;p}$ from $\mathcal{D}$ to arrange that $[T^0]\mathcal{D}=0$ without affecting the fact that $\mathcal{D}$ annihilates the space of theta functions. Then $T^{-1}\mathcal{D} \in \hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, (d-2)s+(d'-2)f)$ again annihilates the space of $BC_1(q\eta)$-symmetric theta functions of degree $l-1$. Thus, by induction, we reduce to the case $d\le l$. By the following lemma, there are no operators of order $<l$ that annihilate all $BC_1(q\eta)$-symmetric theta functions of degree $l-1$, and such an operator of order l is unique up to left multiplication by a meromorphic function. In particular, when $d=l$, we obtain an operator of the form $g\mathcal{D}_{l;\eta;q;p}$ for some meromorphic function g. Comparing the coefficients of T^0 and T^l shows that $g(\eta/q^{l-1}z)=g(z)$, and since $g\mathcal{D}_{l;\eta;q;p} \in \hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ls+d'f)$, we find that f must be holomorphic. □

Lemma 4.18. *Let $f_1, \ldots, f_l$ be a basis of the space of $BC_1(\eta)$-symmetric theta functions of degree $l-1$. Then*

$$\det_{1\le i,j\le l} f_i(z_j) \neq 0 \tag{4.55}$$

unless $z_i \in p^{\mathbb{Z}} z_j \cup p^{\mathbb{Z}}\eta/z_j$ for some $i<j$. In particular, as a function of z,

$$\det_{1\le i,j\le l} f_i(q^{j-1}z) \neq 0 \tag{4.56}$$

unless $q^j \in p^{\mathbb{Z}}$ for some $1 \le j < l$.

Proof. The second claim is an immediate consequence of the first, which follows from the usual degree considerations. To be precise, the given conditions certainly force the determinant to vanish, and thus the determinant is a holomorphic multiple of

$$\prod_{i<j} z_i^{-1}\theta_p(z_i/z_j, z_i z_j/\eta). \tag{4.57}$$

Since this is already a $BC_1(\eta)$-symmetric theta function of degree $l-1$ in each variable, the remaining factor of the determinant is a constant, nonzero since $f_1,\ldots, f_l$ are a basis. □

Corollary 4.19. *If $q^{\mathbb{Z}} \cap p^{\mathbb{Z}} = 1$, then*

$$\dim(\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds+d'f)) = (d+1)(d'+1) \tag{4.58}$$

unless $\eta/\eta' \in q^l p^{\mathbb{Z}}$ for some l with $1 \le l \le \min(d-d', d+1)$, when

$$\dim(\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds+d'f)) = (d-l+1)(d'+l+1). \tag{4.59}$$

Proof. Indeed, we either have

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds+d'f) = \mathcal{S}_{\eta,\eta';q;p}(0, ds+d'f) \tag{4.60}$$

or

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds+d'f) = \mathcal{S}_{\eta,\eta';q;p}(ls-lf, ds+d'f)\mathcal{D}_{l;\eta;q;p}, \tag{4.61}$$

as described. (When $l = d-d'$, both are true.) □

Corollary 4.20. *Let q have order r in $\mathbb{C}^*/\langle p\rangle$, and let $1 \le l \le r$. Then*

$$\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, k(s-f)) = 0 \tag{4.62}$$

unless $k \in r\mathbb{N} \cup l + r\mathbb{N}$, when it has dimension 1. Moreover,

$$\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, ds+d'f) = \hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(k(s-f), ds+d'f)\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0, k(s-f)) \tag{4.63}$$

where

$$k = \begin{cases} 0 & 0 \le d-d' < l \\ mr+l & 2mr+l \le d-d' < (2m+1)r+l, \text{ some } m \ge 0, \\ (m+1)r & (2m+1)r+l \le d-d' < (2m+2)r+l, \text{ some } m \ge 0. \end{cases} \tag{4.64}$$

Proof. We have

$$\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0,ds+d'f)=\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(l(s-f),ds+d'f)\mathcal{D}_{l;\eta;q;p} \tag{4.65}$$

as long as $l\leq \min(d+1,d-d')$. Now,

$$\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(l(s-f),ds+d'f)=\hat{\mathcal{S}}_{q^{-l}\eta,\eta;q;p}(0,(d-l)s+(d'+l)f). \tag{4.66}$$

If $l=r$, then

$$\hat{\mathcal{S}}_{q^{-r}\eta,\eta;q;p}(0,(d-l)s+(d'+l)f).=\hat{\mathcal{S}}_{\eta,\eta;q;p}(0,(d-l)s+(d'+l)f).=\hat{\mathcal{S}}_{\eta,q^{-r}\eta;q;p}(0,(d-l)s+(d'+l)f), \tag{4.67}$$

while if $1\leq l<r$,

$$\hat{\mathcal{S}}_{q^{-l}\eta,\eta;q;p}(0,(d-l)s+(d'+l)f)=\hat{\mathcal{S}}_{q^{-l}\eta,q^{l-r}(q^{-l}\eta);q;p}(0,(d-l)s+(d'+l)f). \tag{4.68}$$

In particular, we may factor out $\mathcal{D}_{r-l;\eta;q;p}$ or $\mathcal{D}_{r;\eta;q;p}$ as appropriate, unless the appropriate inequality holds. The result follows by induction. □

Proposition 4.21. *If q is a primitive r-th root of unity, then for $1\leq l\leq r$, then the* Hom *space $\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0,r(s-f))$ is spanned by the operator*

$$\frac{z^r}{\theta_{p^r}(z^{2r}/\eta^r)}(1-T^r). \tag{4.69}$$

Proof. If $l=r$, this operator and $\mathcal{D}_{r;\eta;q;p}$ have the same leading coefficient and annihilate the space of $BC_1(q\eta)$-symmetric theta functions of degree $r-1$, so are the same operator.

For $1\leq l<r$, we find that for any $BC_1(q\eta)$-symmetric theta function g of degree $r-l$,

$$\frac{z^r}{\theta_{p^r}(z^{2r}/\eta^r)}(1-T^r)g\in\hat{\mathcal{S}}_{\eta,\eta;q;p}(0-(r-l)f,r(s-f))=\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0,rs-lf) \tag{4.70}$$

But T^r commutes with g, so that

$$g\frac{z^r}{\theta_{p^r}(z^{2r}/\eta^r)}(1-T^r)\in\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0,rs-lf) \tag{4.71}$$

for all g, implying that

$$\frac{z^r}{\theta_{p^r}(z^{2r}/\eta^r)}(1-T^r)\in\hat{\mathcal{S}}_{\eta,q^{-l}\eta;q;p}(0,rs-rf) \tag{4.72}$$

as required. □

Now, there exists a sheaf $V_{d;\eta,\eta';q;p}$ on $\mathbb{P}^1$ for each $d\geq 0$ such that

$$\hat{\mathcal{S}}_{\eta,\eta';q;p}(0,ds+d'f)=\Gamma(V_{d;\eta,\eta';q;p}(d')), \tag{4.73}$$

and it will be useful below to understand the structure of this sheaf. It would, in fact, be simple enough to work out the structure explicitly, but for our purposes, the following inductive fact will suffice.

Lemma 4.22. *Let l be the smallest positive integer such that $\eta/\eta' \in p^{\mathbb{Z}}q^l$, or ∞ if no such integer exists. If $l > d$, then*

$$V_{d;\eta,\eta';q;p} \cong \mathcal{O}_{\mathbb{P}^1}^{d+1}, \tag{4.74}$$

and otherwise

$$V_{d;\eta,\eta';q;p} \cong V_{d-l;q^{-l}\eta,q^l\eta';q;p}(l) \oplus \mathcal{O}_{\mathbb{P}^1}(l-d-1)^l. \tag{4.75}$$

Proof. Since $\hat{\mathcal{S}}_{\eta,\eta';q;p}$ is a domain (inherited from the algebra of difference operators), the sheaf $V_{d;\eta,\eta';q;p}$ is torsion-free, so a vector bundle on $\mathbb{P}^1$. Moreover, since we know how many global sections sufficiently large twists of this sheaf have, we may compute its Hilbert polynomial: $h(t; V_{d;\eta,\eta';q;p}) = (d+1)(t+1)$. If $\eta/\eta' \notin p^{\mathbb{Z}}q^l$ for $1 \leq l \leq d$, then $V_{d;\eta,\eta';q;p}(-1)$ has no global sections, so is acyclic, and thus $V_{d;\eta,\eta';q;p} \cong \mathcal{O}_{\mathbb{P}^1}^{d+1}$. Otherwise, let l be the smallest such integer. Then left-multiplication by $\mathcal{D}_{l;\eta;q;p}$ induces the injective map in a short exact sequence

$$0 \to V_{d-l;q^{-l}\eta,q^l\eta';q;p}(l) \to V_{d;\eta,\eta';q;p} \to \mathcal{Q} \to 0. \tag{4.76}$$

We know the Hilbert polynomials of the first two sheaves, and thus find $h(t;\mathcal{Q}) = l(t-d+l)$. Now, $h^0(V_{d;\eta,\eta';q;p}(d-l)) = (d+1)(d-l+1) = \chi(V_{d;\eta,\eta';q;p}(d-l))$, and thus $V_{d;\eta,\eta';q;p}(d-l)$ is acyclic. Moreover, as a quotient of an acyclic sheaf, the same holds for $\mathcal{Q}(d-l)$, and since $\chi(\mathcal{Q}(d-l)) = 0$, we find that $\mathcal{Q}(d-l)$ also has no global sections; it follows that $\mathcal{Q} \cong \mathcal{O}_{P^1}(l-d-1)^l$. Now,

$$V_{d-l;q^{-l}\eta,\eta;q;p}(d-1) \tag{4.77}$$

is acyclic (since $d-1 \geq d-l-1$), and thus the above short exact sequence must split, giving

$$V_{d;\eta,q^{-l}\eta;q;p} \cong V_{d-l;q^{-l}\eta,\eta;q;p}(l) \oplus \mathcal{O}_{P^1}(l-d-1)^l \tag{4.78}$$

as required. □

Of course, each Hom space in $\hat{\mathcal{S}}$ is not only a $\mathbb{P}^1$-module on the left, but also on the right. We may thus view the morphisms in $\hat{\mathcal{S}}$ as global sections of suitable sheaves on $\mathbb{P}^1 \times \mathbb{P}^1$. Each such sheaf has locally free direct image under either projection, and the family of sheaves is closed under twisting by line bundles. It follows that the support of the sheaf must be 1-dimensional, with no component contained in a fiber of either ruling. In other words, $\hat{\mathcal{S}}$ is the space of global sections of a "coherent sheaf bimodule" over $\mathbb{P}^1$ in the terminology of [65], and these fit together to form a sheaf $\mathbb{Z}$-algebra. That is, for each d_1, $d_2 \in \mathbb{Z}$, we have a sheaf S_{d_1,d_2} on $\mathbb{P}^1 \times \mathbb{P}^1$ such that

$$\Gamma(S_{d_1,d_2}(-d_1', d_2')) = \hat{\mathcal{S}}(d_1 s + d_1' f, d_2 s + d_2' f), \tag{4.79}$$

and the composition law induces a morphism

$$\pi_{13*}(\pi_{12}^* S_{d_1,d_2} \otimes \pi_{23}^* S_{d_2,d_3}) \to S_{d_1,d_3}, \tag{4.80}$$

where π_{ij} are the appropriate projection maps from $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1$. Now, Van den Bergh defined the category of sheaves on such a sheaf $\mathbb{Z}$-algebra. When translated back to $\hat{\mathcal{S}}$, we obtain essentially the same definition as above, with one important difference: our definition of torsion above had quantifiers

$$\exists D. \exists D'. \forall d \geq D. \forall d' \geq D' - D + d, \tag{4.81}$$

while Van den Bergh's definition involves the quantifiers

$$\exists D. \forall d \geq D. \exists D'. \forall d' \geq D'. \tag{4.82}$$

(More precisely, Van den Bergh's definition states that for $d \geq D$, every image of v in degree $ds + *f$ is torsion as an element of the corresponding horizontal $\mathbb{Z}$-module.) Of course, it follows immediately from Lemma 4.1 that the two notions are equivalent. In other words, the category of sheaves on $\mathcal{S}_{\eta,\eta';q;p}$ is precisely the same as the category of sheaves on the corresponding sheaf $\mathbb{Z}$-algebra.

Since $\mathcal{S}$ is generated in degrees s, f, $s+f$, it follows that the sheaf $\mathbb{Z}$-algebra is generated in degree 1. The generating sheaves $S_{d,d+1}$ are easily seen to be supported on the curve $\mathbb{C}/\langle p \rangle$, embedded in $\mathbb{P}^1 \times \mathbb{P}^1$ by the pair of involutions $z \mapsto q^{1-d}\eta/z$, $z \mapsto q^{-d}\eta/z$. Moreover, $S_{d,d+1}$ is an invertible sheaf on its support, namely the line bundle corresponding to theta functions with multiplier $q\eta'/pz^2$. The quadratic relations similarly form a sheaf $R_{d,d+2}$ on $\mathbb{P}^1 \times \mathbb{P}^1$ with direct images of rank 1. In this case, the two spaces of theta functions are related by translation by q, so there is a natural way to identify the coordinates. If we do so, then we find that $R_{d,d+2}$ is the structure sheaf of the diagonal. Moreover, this subsheaf of $S_{d,d+1} \otimes S_{d+1,d+2}$ is nondegenerate in the sense of [65]; indeed, otherwise $\hat{\mathcal{S}}$ would have a zero-divisor. Since S_{d_1,d_2} has rank $d_2 - d_1 + 1$, we conclude that the category of sheaves on $\hat{\mathcal{S}}$ is a noncommutative Hirzebruch surface.

We should note in this context that the representation in difference operators is new (and extends to all noncommutative ruled surfaces satisfying a certain separability condition [44]), as are the presence of a canonical family of subcategories $\mathcal{S}$ in which all Hom spaces are flat, and the above formulae for the dimensions of the Hom spaces in $\hat{\mathcal{S}}$.

We can also give a formula for the operators $\mathcal{D}_{l;\eta;q;p}$ in general. We work by first assuming that q is nontorsion, then take a limit to get the torsion case. We can calculate $\mathcal{D}_{l;\eta;q;p}$ in several different ways: e.g., we know the leading coefficient, so can solve for the coefficients using Cramer's rule and the fact that it annihilates all symmetric theta functions of degree $l-1$. Another approach uses the fact that

$$g_1 \cdots g_l \mathcal{D}_{l;\eta;q;p} = \mathcal{D}_{q^{1-l}\eta;q;p}(g_1)\mathcal{D}_{q^{2-l}\eta;q;p}(g_2)\cdots\mathcal{D}_{\eta;q;p}(g_l) = g_1 \cdots g_{l-1}\mathcal{D}_{l-1;\eta/q;q;p}\mathcal{D}_{\eta;q;p}(g_l), \quad (4.83)$$

so that for any $BC_1(q^{1-l}\eta)$-symmetric theta function g of degree 1,

$$g\mathcal{D}_{l;\eta;q;p} = \mathcal{D}_{l-1;\eta/q;q;p}\mathcal{D}_{\eta;q;p}(g). \quad (4.84)$$

If g vanishes at w, we can use this to give a simple first-order recurrence for the values of the coefficients of $\mathcal{D}_{l-1;\eta/q;q;p}$ at w. Either approach gives the following answer.

Proposition 4.23. *For all integers $l \geq 0$,*

$$\mathcal{D}_{l;\eta;q;p} = \sum_{0 \leq k \leq l} \frac{q^{lk}\theta_p(q^{-l};q)_k}{\theta_p(q;q)_k} \frac{z^l \theta_p(q^{2k-1}z^2/\eta)}{\theta_p(q^{k-1}z^2/\eta;q)_{l+1}} T^k. \quad (4.85)$$

Remark. Note that the coefficient

$$\frac{q^{lk}\theta_p(q^{-l};q)_k}{\theta_p(q;q)_k} \quad (4.86)$$

is essentially the elliptic analogue of a q-binomial coefficient. In particular, the zeros of the denominator at torsion q are all cancelled by zeros of the numerator, so the limit to torsion q always exists. Also, in the limit $p \to 0$, we obtain the l-th power of the Askey-Wilson lowering operator, according to a formula of [15]. One should also compare the formulas of [28] and [55]. Note that the latter (up to shifting by a half-integer power of q) can be more directly interpreted as giving the generator of the 1-dimensional Hom space $\mathcal{S}'_{\eta,x_0;q;\mathbb{C}^*/p}(0, ds)$ in the algebraic category $\mathcal{S}'$ constructed below.

It is worth noting one last thing about these operators. Although these operators cause the Fourier transform to fail for $\hat{\mathcal{S}}_{\eta,\eta';q;p}$, they also in a sense *induce* the Fourier transform. Suppose $l > d' - d$ (and q is not torsion, for simplicity), and consider an operator

$$\mathcal{D} \in \hat{\mathcal{S}}_{\eta,q^{d'-d-l}\eta;q;p}(0, ds + d'f) = \mathcal{S}_{\eta,q^{d'-d-l}\eta;q;p}(0, ds + d'f) \tag{4.87}$$

Such an operator takes $BC_1(q\eta)$-symmetric theta functions of degree $d - d' + l - 1$ to functions of degree $l - 1$. It follows, therefore, that $\mathcal{D}_{l;q^{-d}\eta;q;p}\mathcal{D}$ annihilates all $BC_1(q\eta)$-symmetric theta functions of degree $d - d' + l - 1$. As a result, there exists an operator $\hat{\mathcal{D}}$ such that

$$\mathcal{D}_{l;q^{-d}\eta;q;p}\mathcal{D} = \hat{\mathcal{D}}\mathcal{D}_{d-d'+l;\eta;q;p} \tag{4.88}$$

In this way, we obtain an operator in

$$\hat{\mathcal{S}}_{q^{d'-d-l}\eta,\eta;q;p}(0, d's + df) = \mathcal{S}_{q^{d'-d-l}\eta,\eta;q;p}(0, d's + df). \tag{4.89}$$

Since this acts on parameters and objects just as the Fourier transform, we might expect that it agrees on $\mathcal{S}$, and this is indeed the case. It suffices to check this for elements of degree f and s, where we have already seen

$$g\mathcal{D}_{l;\eta;q;p} = \mathcal{D}_{l-1;\eta/q;q;p}\mathcal{D}_{\eta;q;p}(g) \tag{4.90}$$

and taking adjoints easily gives

$$\mathcal{D}_{q^{1-l}\eta;q;p}(g)\mathcal{D}_{l-1;\eta;q;p} = \mathcal{D}_{l;\eta;q;p}g. \tag{4.91}$$

Though this only works for special values of the parameters, we can use the above formula to analytically continue the operators $\mathcal{D}_{l;\eta;q;p}$ to formal operators, and use this to define the Fourier transform whenever q is not torsion. We avoid this approach here, since the existing approach works for all parameters, but this idea will be the key to defining the multivariate analogue of the transform in [46]. It may be worth keeping this approach in mind even in the univariate case, since it gives a way to compute transforms of operators without having to express them in terms of the generators. For more discussion of these formal operators (and in particular the multivariate analogue), see [48, §4].

5 F_1 and $\mathbb{P}^2$

Before we discuss blowing up $\mathcal{S}_{\eta,\eta';q;p}$, it will be useful to consider the analogous deformation of the Hirzebruch surface F_1. Of course, we would want to consider that case in any event, as it is one of the only two kinds of elliptic rational surface that is *not* a blow up of $\mathbb{P}^1 \times \mathbb{P}^1$ or F_2. (The other surface, $\mathbb{P}^2$, will appear as a $\mathbb{Z}$-algebra contained in the F_1 deformation, though as we will see, there are some subtle issues that arise in that case.)

The Picard lattice of F_1 is spanned by elements s and f with $s^2 = -1$ (the unique section of minimal self-intersection) and $f^2 = 0$ (the class of a fiber), with $s \cdot f = 1$. An operator of degree $ds + d'f$ in the corresponding $\mathbb{Z}^2$-algebra should again be a d-th order difference operator, and the formula given for the multiplier of the leading coefficient in the $\mathbb{P}^1 \times \mathbb{P}^1$ case suggests that for F_1 that multiplier should be

$$(-1)^d z^{-d-2d'} q^{d(d+1)/2+d'-dd'} (\eta/p)^{d'} x_0^d \tag{5.1}$$

for a morphism from 0 to $ds + d'f$, and the appropriate ratio more generally. (Of course, the above guess is not particularly unique, but turns out to have the best behavior with respect to elementary transformations.)

This suggests that we should take the following generators:

- If $g(z)$ is a $BC_1(q^{1-d}\eta)$-symmetric theta function of degree 1, then

$$g(z) \in \mathcal{S}'_{\eta,x_0;q;p}(ds + d'f, ds + (d'+1)f). \tag{5.2}$$

- If $h(z)$ is a holomorphic function with $h(pz)/h(z) = -q^{d-d'+1}x_0/z$, then

$$\mathcal{D}_{q^{-d}\eta;q;p}(h) \in \mathcal{S}'_{\eta,x_0;q;p}(ds + d'f, (d+1)s + d'f). \tag{5.3}$$

- If $b(z)$ is a holomorphic function with $b(pz)/b(z) = -q^{1-d'}\eta x_0/pz^3$, then

$$\mathcal{D}_{q^{-d}\eta;q;p}(b) \in \mathcal{S}'_{\eta,x_0;q;p}(ds + d'f, (d+1)s + (d'+1)f). \tag{5.4}$$

We define $\mathcal{S}'_{\eta,x_0;q;p}$ to be the $\mathbb{Z}s + \mathbb{Z}f$-algebra generated by these morphisms (with composition via multiplication of difference operators). Note that this again satisfies an invariance under twisting, in the form

$$\mathcal{S}'_{\eta,x_0;q;p}(d_1s + d'_1f, d_2s + d'_2f) = \mathcal{S}'_{q^{-d_1}\eta,q^{d_1-d'_1}x_0;q;p}(0, (d_2-d_1)s + (d'_2-d'_1)f) \tag{5.5}$$

We should note here that since the generators of degree s depend on a theta function with a single zero, they are unique up to scalar multiplication, and have the form

$$\mathcal{D}_{q^{-d}\eta;q;p}(\theta_p(z/q^{d-d'+1}x_0)). \tag{5.6}$$

The multiples of this operator account for a 2-dimensional subspace of the 3-dimensional space of operators of degree $s+f$, so we again have a set of generators of degree 4 (i.e., each object is the domain of four generators). These difference operators again satisfy commutation relations.

Lemma 5.1. *For any holomorphic function $b(z)$ with $b(pz)/b(z) = -q\eta x_0/pz^3$,*

$$\mathcal{D}_{\eta/q;q;p}(\theta_p(z/qx_0))\mathcal{D}_{\eta;q;p}(b) = \mathcal{D}_{\eta/q;q;p}(b)\mathcal{D}_{\eta;q;p}(\theta_p(z/qx_0)). \tag{5.7}$$

Proof. Again, the coefficients of T^0 and T^2 in the difference vanish, while the coefficient of T is a $BC_1(\eta/q)$-symmetric theta function of degree -2, so also vanishes. □

Also, as in the even case, we can identify the general first-order operator that appears as a morphism in our $\mathbb{Z}^2$-algebra.

Lemma 5.2. *If $b(z)$ is holomorphic with*

$$b(pz)/b(z) = -qx_0(\eta/p)^{d'}z^{-2d'-1}, \tag{5.8}$$

then

$$\mathcal{D}_{\eta;q;p}(b) \in \mathcal{S}'_{\eta,x_0;q;p}(0, s + d'f). \tag{5.9}$$

Proof. For $d' < 0$, there is no such function, while for $d' = 0$ and $d' = 1$, the given operator is a generator. For $d' > 1$, the result follows from Lemma 2.3 by induction. □

The canonical class in F_1 is $-2s-3f$, so we should expect to find the operator T in the negative of that degree.

Lemma 5.3. *We have*

$$T \in \mathcal{S}'_{\eta,x_0;q;p}(ds + d'f, (d+2)s + (d'+3)f). \tag{5.10}$$

Proof. From Lemma 3.3, we have for any $\eta' \in \mathbb{C}^*$ an expansion of the form

$$T = \sum_i \mathcal{D}_{\eta/q;q;p}(b_{i1})\mathcal{D}_{\eta;q;p}(b_{i2}), \tag{5.11}$$

where

$$b_{i1}(pz) = (\eta\eta'/p^2qz^4)b_{i1}(z),$$
$$b_{i2}(pz) = (q\eta\eta'/p^2z^4)b_{i2}(z),$$

and $b_{i1}(q\eta'/x_0) = 0$ for all i. Suppose $|q| < 1$, and write

$$\begin{aligned} T &= \Gamma_{p,q}(q^2z\eta'/\eta x_0, q\eta'/x_0z)T\Gamma_{p,q}(qz\eta'/\eta x_0, q^2\eta'/x_0z)^{-1} \\ &= \sum_i \Gamma_{p,q}(q^2z\eta'/\eta x_0, q\eta'/x_0z)\mathcal{D}_{\eta/q;q;p}(b_{i1})\Gamma_{p,q}(q^2z\eta'/\eta x_0, q^2\eta'/x_0z)^{-1} \\ &\qquad \Gamma_{p,q}(q^2z\eta'/\eta x_0, q^2\eta'/x_0z)\mathcal{D}_{\eta;q;p}(b_{i2})\Gamma_{p,q}(qz\eta'/\eta x_0, q^2\eta'/x_0z)^{-1} \\ &= \sum_i \mathcal{D}_{\eta/q;q;p}(b_{i1}\theta_p(q\eta'/x_0z)^{-1})\mathcal{D}_{\eta;q;p}(b_{i2}\theta_p(qz\eta'/\eta x_0)). \end{aligned} \tag{5.12}$$

Since this expansion is an algebraic identity, it continues to hold for general q. The result follows by observing that the first factor of each summand is an element of $\mathcal{S}'_{\eta,x_0;q;p}(s+2f, 2s+3f)$, while the second factor is an element of $\mathcal{S}'_{\eta,x_0;q;p}(0, s+2f)$. □

We can then prove flatness as in the even case; the main difference is that now the presentation is *always*, rather than merely generically valid.

Theorem 5.4. *For any integers d, d',*

$$\dim(\mathcal{S}'_{\eta,x_0;q;p}(0, ds+d'f)) = \begin{cases} 0 & d < 0 \\ (d+1)(2d'+2-d)/2 & 0 \le d \le d' \\ (d'+1)(d'+2)/2 & d \ge d'. \end{cases} \tag{5.13}$$

In particular, $\mathcal{S}'_{\eta,x_0;q;p}$ is a flat family of $\mathbb{Z}^2$-algebras.

Proof. The dimension is certainly 0 if $d < 0$. If $d > d'$, then we have

$$\mathcal{S}'_{\eta,x_0;q;p}(0, ds+d'f) \supset \mathcal{S}'_{\eta,x_0;q;p}((d-1)s+d'f, ds+d'f)\mathcal{S}'_{\eta,x_0;q;p}(0, (d-1)s+d'f), \tag{5.14}$$

so that

$$\dim(\mathcal{S}'_{\eta,x_0;q;p}(0, ds+d'f)) \ge \dim(\mathcal{S}'_{\eta,x_0;q;p}(0, (d-1)s+d'f)). \tag{5.15}$$

If $d < d'$, then we find that all functions with the appropriate multiplier and divisor occur as leading coefficients, and thus

$$\dim(\mathcal{S}'_{\eta,x_0;q;p}(0, ds+d'f)) \ge d+2d' + \dim(\mathcal{S}'_{\eta,x_0;q;p}(0, (d-2)s+(d'-3)f)). \tag{5.16}$$

If $d = d'$, we may combine the two inequalities to find

$$\dim(\mathcal{S}'_{\eta,x_0;q;p}(0, d(s+f))) \ge 3d + \dim(\mathcal{S}'_{\eta,x_0;q;p}(0, (d-3)(s+f)). \tag{5.17}$$

We thus find by induction that the claimed formula gives a lower bound on the dimension.

To prove the corresponding upper bound, we choose four generators as described above, and note that using quadratic relations, we may express any element of $\mathcal{S}'_{\eta,x_0;q;p}(0, ds + d'f)$ in the form

$$\sum_k \sum_{0 \le j \le d'-k} \alpha_{kj} x_0^{d'-k-j} x_1^j y_0^{d-k} y_1^k, \tag{5.18}$$

where x_0, x_1 are the two generators of degree f, and y_0, y_1 are the generators of degree s, $s+f$ respectively. This gives the upper bound

$$\dim(\mathcal{S}'_{\eta,x_0;q;p}(0, ds + d'f)) \ge \sum_{0 \le k \le \min(d,d')} (d' - k + 1), \tag{5.19}$$

agreeing with the above formula. □

Corollary 5.5. *If $d > d' \ge 0$, then*

$$\begin{aligned} \mathcal{S}'_{\eta,x_0;q;p}(0, ds + d'f) &= \mathcal{S}'_{\eta,x_0;q;p}(d's + d'f, ds + d'f)\mathcal{S}'_{\eta,x_0;q;p}(0, d's + d'f) \\ &= \mathcal{S}'_{\eta,x_0;q;p}((d-d')s, ds + d'f)\mathcal{S}'_{\eta,x_0;q;p}(0, (d-d')s). \end{aligned} \tag{5.20}$$

Remark. This should be compared to the fact that on a commutative surface, if the divisor class D has negative intersection with a -1-curve, then it has that -1-curve as a fixed component with multiplicity the negative of the intersection. (Indeed, the intersection pairing on F_1 has $s \cdot (ds + d'f) = d' - d$.) Note that the behavior of -1-curves in the noncommutative setting is thus much simpler than that of -2-curves.

Analogously to the even case, we may define sheaves on $\mathcal{S}'$ by calling an element torsion if for some $d' \ge d \ge 0$ it is annihilated by all elements of degree $(d + d_1)s + (d' + d'_1)f$ for any $d'_1 \ge d_1 \ge 0$. (Again, this is the cone of divisors which are nef on all commutative surfaces in the family.) The same arguments carry over to the odd case, so in particular we find that both the natural representation in symmetric meromorphic functions and the regular representation are already saturated.

Just as in the F_0/F_2 case, we could have obtained the corresponding category of sheaves from the construction of [65]. Indeed, we find that a homogeneous element v of a module is torsion iff there exists d_0 such that for $d \ge d_0$, any image of v under a morphism of degree $ds + *f$ is torsion w.r.to morphisms of degree a multiple of f (i.e., is 0 as an element of the appropriate sheaf on $\mathbb{P}^1$). This exhibits our category of sheaves as the category of sheaves associated to a sheaf $\mathbb{Z}$-algebra on $\mathbb{P}^1$. As in the F_0/F_2 case, we find that the resulting noncommutative scheme is a noncommutative Hirzebruch surface.

The odd case has another useful special feature. Since the morphisms of degree s are unique up to scalar multiplication, they give rise to a well-defined family of difference equations. It turns out that we can use the solutions to those equations to produce a canonical version of the $\mathbb{Z}^2$-algebra in which the morphisms are operators with *elliptic* (as opposed to meromorphic theta function) coefficients. To be precise, if $|q| < 1$, the function

$$\Gamma_{p,q}(z/q^{d+1-d'}x_0, \eta/q^{2d-d'}x_0 z) \tag{5.21}$$

is annihilated by all sections of $\mathcal{S}'_{\eta,x_0;q;p}(ds + d'f, (d+1)s + d'f)$, and has the following property.

Proposition 5.6. *Suppose $|q| < 1$. For any operator*

$$\mathcal{D} \in \mathcal{S}'_{\eta,x_0;q;p}(d_1 s + d'_1 f, d_2 s + d'_2 f), \tag{5.22}$$

the coefficients of the operator

$$\Gamma_{p,q}(z/q^{d_2+1-d_2'}x_0,\eta/q^{2d_2-d_2'}x_0z)^{-1}\mathcal{D}\Gamma_{p,q}(z/q^{d_1+1-d_1'}x_0,\eta/q^{2d_1-d_1'}x_0z) \tag{5.23}$$

are invariant under $z\mapsto pz$.

Proof. This transformation of the operators respects composition, as does the property of having elliptic coefficients. It thus remains only to show that the generators transform to elliptic difference operators. For degree f, we have

$$\begin{aligned}&\Gamma_{p,q}(z/q^{d-d'}x_0,\eta/q^{2d-d'-1}x_0z)^{-1}g(z)\Gamma_{p,q}(z/q^{d+1-d'}x_0,\eta/q^{2d-d'}x_0z)\\&=\frac{g(z)}{\theta_p(z/q^{d+1-d'}x_0,\eta/q^{2d-d'}x_0z)};\end{aligned} \tag{5.24}$$

for degree s, we have

$$\begin{aligned}&\Gamma_{p,q}(z/q^{d+2-d'}x_0,\eta/q^{2d+2-d'}x_0z)^{-1}\mathcal{D}_{\eta;q;p}(\theta_p(z/q^{d-d'+1}x_0))\Gamma_{p,q}(z/q^{d+1-d'}x_0,\eta/q^{2d-d'}x_0z)\\&=\frac{z\theta_p(q^{d'-d-1}z/x_0,q^{d'-d-2}z/x_0,q^{d'-2d-1}\eta/x_0z,q^{d'-2d-2}\eta/x_0z)}{\theta_p(q^dz^2/\eta)}(1-T);\end{aligned} \tag{5.25}$$

and for degree $s+f$ we have

$$\begin{aligned}&\Gamma_{p,q}(z/q^{d+1-d'}x_0,\eta/q^{2d+1-d'}x_0z)^{-1}\mathcal{D}_{\eta;q;p}(b)\Gamma_{p,q}(z/q^{d+1-d'}x_0,\eta/q^{2d-d'}x_0z)\\&=\mathcal{D}_{\eta;q;p}(b(z)\theta_p(q^{d'-2d-1}\eta/x_0z)).\end{aligned} \tag{5.26}$$

In each case, we can directly check that the coefficients are elliptic. □

This allows us to define an analogue of $\mathcal{S}'$ on an *algebraic* elliptic curve. That is, if C is a smooth genus 1 curve (now over an arbitrary field, possibly of finite characteristic) and we are given parameters $\eta\in\mathrm{Pic}^2(C)$, $x_0\in\mathrm{Pic}^1(C)$, $q\in\mathrm{Pic}^0(C)$, then we can define a $\mathbb{Z}^2$-algebra $\mathcal{S}'_{\eta,x_0;q;C}$ generated by the following elliptic q-difference operators. Given $\eta\in\mathrm{Pic}^2(C)$, we define R_η to be the corresponding ramification divisor; i.e., R_η is the divisor associated to the subscheme of fixed points of the involution $z\mapsto\eta/z$.

- If $g\in\mathcal{L}([q^{d+1-d'}x_0][\eta/q^{2d-d'}x_0])$, then $g\in\mathcal{S}'_{\eta,x_0;q;C}(ds+d'f,ds+(d'+1)f)$.
- If $h\in\mathcal{L}(R_{q^{-d}\eta}/[q^{d+1-d'}x_0][q^{d+2-d'}x_0][\eta/q^{2d+1-d'}x_0][\eta/q^{2d+2-d'}x_0])$, then

$$h(z)+h(q^{-d}\eta/z)T=h(z)(1-T)\in\mathcal{S}'_{\eta,x_0;q;C}(ds+d'f,(d+1)s+d'f). \tag{5.27}$$

- If $b\in\mathcal{L}(R_{q^{-d}\eta}/[\eta/q^{2d+1-d'}x_0])$, then

$$b(z)+b(q^{-d}\eta/z)T\in\mathcal{S}'_{\eta,x_0;q;C}(ds+d'f,(d+1)s+(d'+1)f). \tag{5.28}$$

Remark. Note that in general, we have

$$[T^0]\mathcal{S}_{\eta,x_0;q;C}(0,ds+d'f)\subset\Gamma(C;\mathcal{L}(\prod_{0\le k<d}R_{q^{-k}\eta}\prod_{0\le k<d'-d}[q^{1-k}x_0]\prod_{0\le k<d'-2d}[\eta/q^{-k}x_0])), \tag{5.29}$$

extended in the natural way if $d'<d$ or $d'<2d$; more generally, the leading coefficient of an element of $\mathcal{S}_{\eta,x_0;q;C}(d_1s+d_1'f,d_2s,d_2'f)$ is naturally a section of the line bundle associated to the corresponding ratio of divisors.

The above argument carries over directly; the only nontrivial issue is the commutation relations, which follow by taking a limit.

Corollary 5.7. *This family of categories is flat.*

The above description is clearly functorial in the quadruple (C, η, x_0, q), immediately implying that $\mathcal{S}'_{\eta,x_0;q;p}$ is modular. Two important special cases of functoriality are the isomorphisms

$$\begin{aligned}\mathcal{S}'_{\eta,x_0;q;C} &\cong \mathcal{S}'_{\tau^2\eta,\tau x_0;q;C} \\ &\cong \mathcal{S}'_{\rho^2/\eta,\rho/x_0;1/q;C}\end{aligned}$$

for any $\tau \in \mathrm{Pic}^0(C)$, $\rho \in \mathrm{Pic}^2(C)$; both isomorphisms act as the identity on objects. We also have an analogue of twist-invariance:

$$\mathcal{S}'_{\eta,x_0;q;C} \cong \mathcal{S}'_{q^{-d}\eta,q^{d-d'}x_0;q;C}, \qquad d_0 s + d'_0 f \mapsto (d+d_0)s + (d'+d'_0)f, \tag{5.30}$$

as well as an analogue of the formal adjoint:

$$\mathcal{S}'_{\eta,x_0;q;C} \cong \mathcal{S}'_{\eta,x_0,1/q;C} \qquad d_0 s + d'_0 f \mapsto 2s + 3f - (d_0 s + d'_0 f). \tag{5.31}$$

Analytically, the latter is formally an adjoint with respect to the family of (elliptic) densities

$$\frac{z\theta_p(z/q^{d-d'}x_0, z/q^{d+1-d'}x_0, \eta/q^{2d-d'-1}x_0 z, \eta/q^{2d-d'}x_0 z)}{\theta_p(q^{d-1}z^2/\eta)}. \tag{5.32}$$

Algebraically, the same holds, except that we now need to choose for each $ds + d'f$ an elliptic function with the appropriate divisor (or, equivalently, a generator of $\mathcal{S}'((d-1)s + d'f, ds + d'f)$). We should note in particular that the adjoint involution is thus only defined up to a scalar gauge automorphism. This is mostly unavoidable, but we note that if we choose the scalars so that $\Delta_{(d+2)s+(d'+3)f}(z) = \Delta_{ds+d'f}(qz)$, then the adjoint of T will be T in every degree.

We will have use for the following cancellation lemma in the sequel.

Lemma 5.8. *Let $b_l(z)$ be a holomorphic function such that*

$$b_l(pz)/b_l(z) = -(q^2\eta/pz^2)^l x_0/z, \tag{5.33}$$

and such that if $b_l(z)/\theta_p(z/u_1, z/u_2)$ is holomorphic, then $q\eta/u_1u_2 \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$. Then for any meromorphic difference operator $\mathcal{D}$ of order d, we have

$$\mathcal{D} \in \mathcal{S}'_{\eta,x_0;q;p}(0, ds + d'f) \tag{5.34}$$

iff

$$\mathcal{D}\mathcal{D}_{q\eta;q;p}(b_l) \in \mathcal{S}'_{\eta,x_0;q;p}(-s - lf, ds + d'f). \tag{5.35}$$

Proof. We certainly have $\mathcal{D}_{q\eta;q;p}(b) \in \mathcal{S}'_{\eta,x_0;q;p}(-s - lf, 0)$, so one direction is immediate. For the other direction, we note first that since $\mathcal{S}'_{\eta,x_0;q;p}$ is saturated, it suffices to show that all left-multiples of $\mathcal{D}$ by symmetric theta functions of sufficiently large degree are in $\mathcal{S}'_{\eta,x_0;q;p}$. In particular, we may feel free to assume $d' \geq d$.

Now, let F_0 be a formal symbol on which meromorphic difference operators act via

$$TF_0 = \frac{b_l(z)}{b_l(q\eta/z)}F_0. \tag{5.36}$$

(For $|p|, |q| < 1$, we can take F_0 to be a suitable product of elliptic Gamma functions.) This symbol is annihilated by $\mathcal{D}_{q\eta;q;p}(b_l)$, and thus by the composition $\mathcal{D}\mathcal{D}_{q\eta;q;p}(b_l)$. It will thus suffice to show that any element of $\mathcal{S}'_{\eta,x_0;q;p}(-s-lf, ds+d'f)$ that annihilates F_0 is in the image of $\mathcal{S}'_{\eta,x_0;q;p}(0, ds+d'f)$ under right-multiplication by $\mathcal{D}_{q\eta;q;p}(b_l)$. Equivalently, we need to show that the space

$$\mathcal{S}'_{\eta,x_0;q;p}(-s-lf, ds+d'f)F_0 \tag{5.37}$$

has dimension

$$\dim \mathcal{S}'_{\eta,x_0;q;p}(-s-lf, ds+d'f) - \dim \mathcal{S}'_{\eta,x_0;q;p}(0, ds+d'f) = (l-1)d + 2l + d' \tag{5.38}$$

Now, if we define

$$F_k = \prod_{0\le i<k} b_l(\eta/q^{i-1}z)^{-1}F_0, \tag{5.39}$$

then we find

$$\mathcal{D}_{q^{1-k}\eta;q;p}(b(z))F_k = \frac{b(z)b_l(\eta/q^{k-1}z) - b(\eta/q^{k-1}z)b_l(z)}{z^{-1}\theta_p(q^{k-1}z^2/\eta)}F_{k+1}. \tag{5.40}$$

If $b(z)$ is holomorphic with

$$b(pz)/b(z) = (\eta/pq^{k-1}z^2)^j(-q^{(k+1)l}x_0/z), \tag{5.41}$$

then the coefficient here is a $BC_1(q^{1-k}\eta)$-symmetric theta function of degree $l+j-1$. It follows by an easy induction that

$$\mathcal{S}'_{\eta,x_0;q;p}(-s-lf, ds+d'f)F_0 \tag{5.42}$$

is contained in

$$\mathcal{S}'_{\eta,x_0;q;p}(-s-lf+(d+1)(s-(l-1)f), ds+d'f)F_{d+1}. \tag{5.43}$$

The latter has exactly the required dimension, so we need to prove that the map is surjective; again an inductive argument reduces to the case $d=0$. This then reduces to understanding when

$$b(z)b_l(q\eta/z) - b(q\eta/z)b_l(z) = 0. \tag{5.44}$$

The condition on the multiplier of b then implies that

$$b(z)b_l(q\eta/z) \tag{5.45}$$

is a $BC_1(q\eta)$-symmetric theta function of degree $j+l+1$. Since b_l is not a multiple of any $BC_1(q\eta)$-symmetric theta function, it follows that

$$\frac{b(z)}{b_l(z)} \tag{5.46}$$

must itself be a $BC_1(q\eta)$-symmetric theta function. In particular, it follows that the map

$$\mathcal{S}'_{\eta,x_0;q;p}(-s-lf, -f) \to \mathcal{S}'_{\eta,x_0;q;p}(-(2l-1)f, -f) \tag{5.47}$$

(the case $d=0$, $d'=-1$ above) is injective. For $l>0$, both spaces have dimension $2l-1$, and the map is thus an isomorphism, implying surjectivity for all $d' \ge -1$. For $l=0$, both spaces are trivial, but in that case it is easy to show that

$$\mathcal{S}'_{\eta,x_0;q;p}(-s, f)F_0 = \mathbb{C}F_1, \tag{5.48}$$

again giving surjectivity. □

In the commutative case, F_1 is the blow up of $\mathbb{P}^2$ in one point. We should thus be able to recover a noncommutative deformation of $\mathbb{P}^2$ by "blowing down" s. In the commutative setting, if X is a blowup of Y in a single point, then $\mathrm{Pic}(Y) \subset \mathrm{Pic}(X)$ can be identified as the orthogonal complement of the exceptional curve, and this identifies the natural category with objects $\mathrm{Pic}(Y)$ as a subcategory of the corresponding category for X. In our case, we should thus obtain a noncommutative $\mathbb{P}^2$ by restricting $\mathcal{S}'_{\eta,x_0;q;C}$ to the objects $\mathbb{Z}(s+f)$. And, indeed, the dimension formula above agrees with the Hilbert function of $\mathbb{P}^2$ as expected.

Here, though, we encounter a slightly odd effect. The spaces of operators we obtain clearly depend on η, x_0 and q, and we thus obtain a $2+1$-parameter family. On the other hand, the moduli space of elliptic curves in $\mathbb{P}^2$ is only 1-dimensional, so we should only have a 2-parameter family of deformations. This is indeed the case, but relates to a somewhat unexpected isomorphism.

Lemma 5.9. *The $\mathbb{Z}$-algebra $\mathcal{S}'_{\eta,x_0;q;C}|_{\mathbb{Z}(s+f)}$ is generated in degree $s+f$.*

Proof. We need to show

$$\mathcal{S}'_{\eta,x_0;q;C}(0,(d+1)(s+f)) = \mathcal{S}'_{\eta,x_0;q;C}(d(s+f),(d+1)(s+f))\mathcal{S}'_{\eta,x_0;q;C}(0,d(s+f)). \tag{5.49}$$

Since the generators of $\mathcal{S}'$ have degrees f, s, and $s+f$, we thus reduce to showing

$$\begin{aligned}\mathcal{S}'_{\eta,x_0;q;C}((d+1)s+df,(d+1)(s+f))&\mathcal{S}'_{\eta,x_0;q;C}(0,(d+1)s+df)\\ &\subset \mathcal{S}'_{\eta,x_0;q;C}(d(s+f),(d+1)(s+f))\mathcal{S}'_{\eta,x_0;q;C}(0,d(s+f))\end{aligned} \tag{5.50}$$

and

$$\begin{aligned}\mathcal{S}'_{\eta,x_0;q;C}(ds+(d+1)f,(d+1)(s+f))&\mathcal{S}'_{\eta,x_0;q;C}(0,ds+(d+1)f)\\ &\subset \mathcal{S}'_{\eta,x_0;q;C}(d(s+f),(d+1)(s+f))\mathcal{S}'_{\eta,x_0;q;C}(0,d(s+f)).\end{aligned} \tag{5.51}$$

For the first claim, we note that $s\cdot((d+1)s+df)<0$, and thus

$$\begin{aligned}&\mathcal{S}'_{\eta,x_0;q;C}((d+1)s+df,(d+1)(s+f))\mathcal{S}'_{\eta,x_0;q;C}(0,(d+1)s+df)\\ &=\mathcal{S}'_{\eta,x_0;q;C}((d+1)s+df,(d+1)(s+f))\mathcal{S}'_{\eta,x_0;q;C}(ds+df,(d+1)s+df)\mathcal{S}'_{\eta,x_0;q;C}(0,ds+df)\\ &\subset \mathcal{S}'_{\eta,x_0;q;C}(ds+df,(d+1)(s+f))\mathcal{S}'_{\eta,x_0;q;C}(0,ds+df).\end{aligned} \tag{5.52}$$

For the second claim, we note that the proof of flatness shows that

$$\mathcal{S}'_{\eta,x_0;q;C}(0,ds+(d+1)f) = \mathcal{S}'_{\eta,x_0;q;C}(ds+df,ds+(d+1)f)\mathcal{S}'_{\eta,x_0;q;C}(0,ds+df), \tag{5.53}$$

and the claim again follows immediately. □

In particular, we find that we obtain a $\mathbb{Z}$-algebra in which the generators can be identified with global sections of appropriate line bundles of degree 3, and those generators satisfy the same quadratic relations as those global sections. This presentation appeared in [3], where it was shown that the result was a flat deformation of $\mathbb{P}^2$. It follows in particular that our $\mathbb{Z}$-algebra is quadratic, and isomorphic to that constructed in [11], the $\mathbb{Z}$-algebra analogue of the usual three-generator Sklyanin algebra. (Note that in sharp contrast to the quadratic relations for $\mathcal{S}$ and $\mathcal{S}'$, these relations do not in general form a Gröbner basis.)

In particular, we find that $\mathcal{S}'_{\eta,x_0;q;C}|_{\mathbb{Z}(s+f)}$ depends up to isomorphism only on q and the product bundle ηx_0 (a fact we will give a direct proof of below), but the resulting $\mathbb{Z}$-algebra has a 1-parameter family of difference operator representations. We should note here that this does *not*

give rise in any obvious way to a difference operator representation of the three-generator Sklyanin *algebra*. The difficulty here is that the "central" element (the unique operator of degree $3d+3f$ with vanishing leading coefficient) starts as a *third*-order operator, and thus remains a nontrivial difference operator even when its extreme coefficients vanish. In particular, we cannot in general divide by this operator and still hope to obtain a difference operator. With care, we can divide by the operator *formally*, and even, given a choice of cube root of q, construct a suitable cube root of the operator; the result is a representation of the three-generator Sklyanin algebra in *formal* difference operators. We postpone discussion of this to the multivariate case [46].

One should also note that, as one might expect, the category of sheaves on our noncommutative F_1 is indeed the same category as would be obtained from a noncommutative blowup of the above noncommutative $\mathbb{P}^2$. Indeed, it is straightforward to show that a homogeneous element v is torsion iff there exists r_0 such that for $r > r_0$, any image of v under a morphism of degree $d(s+f)-rs$ is torsion w.r.to the restriction to the corresponding coset of $\mathbb{Z}(s+f)$. This exhibits our category as the category of sheaves on a sheaf $\mathbb{Z}$-algebra on the relevant noncommutative $\mathbb{P}^2$ (which is essentially independent of r). We furthermore see that the sheaf bimodule corresponding to $\mathbb{Z}(s+f)-s$ is precisely the two-sided ideal of a point, and since the sheaf algebra is generated in degree 1, our category is the category of sheaves on a noncommutative blowup (in the sense of [64]) of the appropriate noncommutative $\mathbb{P}^2$. Note that Van den Bergh's blowup construction requires that the point be contained in a commutative curve which is a "divisor", but this follows from the classification of point modules over 3-generator Sklyanin algebras [4]. (Specifically, in our case, the commutative curve ($\cong C$) is cut out by the two-sided ideal $\ker[T^0]$, which is (locally) principal: if

$$\mathcal{D} \in \ker[T^0] \cap \mathcal{S}'_{\eta,x_0;q;C}(d_1(s+f), d_2(s+f)), \tag{5.54}$$

then there is a factorization $\mathcal{D} = T\mathcal{D}_1\mathcal{D}_2$ with $\mathcal{D}_2 \in \mathcal{S}'_{\eta,x_0;q;C}(d_1(s+f),(d_2-3)(s+f))$ and $\mathcal{D}_1$ the unique (mod scalars) section of $\mathcal{S}'_{\eta,x_0;q;C}((d_2-3)(s+f),(d_2-3)(s+f)+s)$, as well as a corresponding factorization on the right.)

6 Surfaces with $K_X^2 = 7$

While we could extend the above definition via generators to one-point blowups of the above Hirzebruch surfaces, it will be more convenient to construct the blowups by imposing appropriate conditions on the morphisms. We will essentially define the blowup as the largest $\mathbb{Z}^3$-algebra that blows down to both $\mathbb{P}^1 \times \mathbb{P}^1$ and F_1.

It turns out that the elementary transformation is relatively straightforward at the level of difference operators, with one mild technicality. Roughly speaking, the transformation involves gauging by a suitable product of elliptic Gamma functions:

$$\Gamma_{p,q}(q^r z/x, q^{r-d+1}\eta/xz). \tag{6.1}$$

Of course, this only makes sense per se when $|q| < 1$; though we could always arrange this by multiplying q by a sufficiently large power of p, with a little care, we can make sense of the gauge transformation regardless. To do this, define a system of *formal* symbols

$$\Gamma_{r;q;p}(z;x) \tag{6.2}$$

with an action of T, such that

$$T\Gamma_{r;q;p}(z;x) = \Gamma_{r+1;q;p}(z;x) = \theta_p(q^r z/x)\Gamma_{r;q;p}(z;x). \tag{6.3}$$

If $|q| < 1$, this can be instantiated by the meromorphic functions $\Gamma_{p,q}(q^r z/x)$, or more precisely by the product of any such function by a q-elliptic function. In general, though, we should think of $\Gamma_{r;q;p}(z;x)$ as a family of explicitly cohomologous 1-cocycles. We thus have well-defined operations

$$\mathcal{D} \mapsto \Gamma_{r_2;q;p}(z;x)^{-1}\mathcal{D}\Gamma_{r_1;q;p}(z;x) \tag{6.4}$$

with the property that

$$(\Gamma_{r_3;q;p}(z;x)^{-1}\mathcal{D}_2\Gamma_{r_2;q;p}(z;x))(\Gamma_{r_2;q;p}(z;x)^{-1}\mathcal{D}_1\Gamma_{r_1;q;p}(z;x)) = \Gamma_{r_3;q;p}(z;x)^{-1}\mathcal{D}_2\mathcal{D}_1\Gamma_{r_1;q;p}(z;x). \tag{6.5}$$

If

$$\mathcal{D} = \sum_{0\le k\le d} c_k(z)T^k, \tag{6.6}$$

then

$$\Gamma_{r_2;q;p}(z;x)^{-1}\mathcal{D}\Gamma_{r_1;q;p}(z;x) = \sum_{0\le k\le d} c_k(z)\theta_p(q^{r_2}z/x;q)_{r_1+k-r_2}T^k, \tag{6.7}$$

where we recall that

$$\theta_p(z;q)_r = \begin{cases} \prod_{0\le j<r}\theta_p(q^j z) & r \ge 0 \\ \prod_{-r\le j<0}\theta_p(q^j z)^{-1} & r \le 0 \end{cases} \tag{6.8}$$

Note in particular that

$$\Gamma_{r_2;q;p}(z;q^{-r}x)^{-1}\mathcal{D}\Gamma_{r_1;q;p}(z;q^{-r}x) = \Gamma_{r_2+r;q;p}(z;x)^{-1}\mathcal{D}\Gamma_{r_1+r;q;p}(z;x). \tag{6.9}$$

Of course, these formal gauge transformations commute with each other, so we need not specify the order in which we perform such operations.

In particular, we may represent the function

$$\Gamma_{p,q}(q^r z/x, q^{r-d+1}\eta/xz) = \frac{\Gamma_{p,q}(q^r z/x)}{\Gamma_{p,q}(q^{d-r}pxz/\eta)} \sim \frac{\Gamma_{r;q;p}(z;x)}{\Gamma_{d-r;q;p}(z;\eta/px)}. \tag{6.10}$$

We should note, of course, that though it is somewhat obscured by the notation, this is still a cocycle for the infinite dihedral group, corresponding to the fact that the original product of Gamma functions is invariant under $z \mapsto q^{1-d}\eta/z$. We also note that Proposition 5.6 above extends naturally to all q if we replace the Gamma functions as above; we find that

$$\frac{\Gamma_{2d_2+1-d_2';q;p}(z;\eta/px_0)}{\Gamma_{d_2'-d_2-1;q;p}(z;x_0)}\mathcal{D}\frac{\Gamma_{d_1'-d_1-1;q;p}(z;x_0)}{\Gamma_{2d_1+1-d_1';q;p}(z;\eta/px_0)} \tag{6.11}$$

is an elliptic difference operator whenever

$$\mathcal{D} \in \mathcal{S}'_{\eta,x_0;q;p}(d_1 s + d_1' f, d_2 s + d_2' f). \tag{6.12}$$

Definition 6.1. $\mathcal{S}_{\eta,\eta',x_1;q;p}$ is the category with objects $\mathbb{Z}s + \mathbb{Z}f + \mathbb{Z}e_1$ and Hom spaces

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(d_1 s + d_1' f - r_{11}e_1, d_2 s + d_2' f - r_{21}e_1) \subseteq \mathcal{S}_{\eta,\eta';q;p}(d_1 s + d_1' f, d_2 s + d_2' f) \tag{6.13}$$

consisting of those operators $\mathcal{D}$ such that

$$\frac{\Gamma_{d_2-r_{21};q;p}(z;\eta/px_1)}{\Gamma_{r_{21};q;p}(z;x_1)}\mathcal{D}\frac{\Gamma_{r_{11};q;p}(z;x_1)}{\Gamma_{d_1-r_{11};q;p}(z;\eta/px_1)} \in \mathcal{S}'_{\eta,\eta'/px_1;q;p}(d_1 s + (d_1 + d_1' - r_{11})f, d_2 s + (d_2 + d_2' - r_{21})f). \tag{6.14}$$

Similarly, $\mathcal{S}'_{\eta,x_0,x_1;q;p}$ is the category with objects $\mathbb{Z}s+\mathbb{Z}f+\mathbb{Z}e_1$ and Hom spaces

$$\mathcal{S}'_{\eta,x_0,x_1;q;p}(d_1s+d'_1f-r_{11}e_1,d_2s+d'_2f-r_{21}e_1)\subseteq\mathcal{S}'_{\eta,x_0;q;p}(d_1s+d'_1f,d_2s+d'_2f) \tag{6.15}$$

consisting of those operators $\mathcal{D}$ such that

$$\begin{aligned}\frac{\Gamma_{d_2-r_{21};q;p}(z;\eta/px_1)}{\Gamma_{r_{21};q;p}(z;x_1)}\mathcal{D}\frac{\Gamma_{r_{11};q;p}(z;x_1)}{\Gamma_{d_1-r_{11};q;p}(z;\eta/px_1)}&\\ \in\mathcal{S}_{\eta,x_0\eta/x_1}(d_1s+(d'_1-r_{11})f,d_2s+(d'_2-r_{21})f)&\end{aligned} \tag{6.16}$$

Note that each Hom space in either category is an intersection of a Hom space and a gauge transformation of a Hom space, and thus the given conditions are indeed closed under composition of difference operators.

The definition clearly inherits a twisting symmetry from $\mathcal{S}_{\eta,\eta';q;p}$ and $\mathcal{S}'_{\eta,x_0;q;p}$, giving the following.

Lemma 6.1. *There are isomorphisms*

$$\begin{aligned}\mathcal{S}_{\eta,\eta',x_1;q;p}&\cong\mathcal{S}_{q^{-d}\eta,q^{-d'}\eta',q^{-r_1}x_1;q;p}\\ \mathcal{S}'_{\eta,x_0,x_1;q;p}&\cong\mathcal{S}'_{q^{-d}\eta,q^{d-d'}x_0,q^{-r_1}x_1;q;p}\end{aligned}$$

which act on objects as $v\mapsto v-ds-d'f+r_1e_1$, and on morphisms as the identity.

This, as usual, allows us to mostly restrict our attention to morphisms from the origin. Note that if

$$\mathcal{D}=\sum_{0\le k\le d}c_k(z)T^k, \tag{6.17}$$

then

$$\frac{\Gamma_{d-r_1;q;p}(z;\eta/px_1)}{\Gamma_{r_1;q;p}(z;x_1)}\mathcal{D}\frac{\Gamma_{0;q;p}(z;x_1)}{\Gamma_{0;q;p}(z;\eta/px_1)}=\sum_{0\le k\le d}\frac{c_k(z)}{\theta_p(q^kz/x_1;q)_{r_1-k}\theta_p(\eta/x_1zq^{k-1};q)_{r_1-d+k}}T^k. \tag{6.18}$$

Lemma 6.2. *We have*

$$\begin{aligned}\frac{\Gamma_{d-r_1;q;p}(z;\eta/px_1)}{\Gamma_{r_1;q;p}(z;x_1)}\mathcal{S}_{\eta,\eta',x_1;q;p}(0,ds+d'f-r_1e_1)\frac{\Gamma_{0;q;p}(z;x_1)}{\Gamma_{0;q;p}(z;\eta/px_1)}&\\ =\mathcal{S}'_{\eta,\eta'/px_1,\eta/px_1;q;p}(0,ds+(d+d'-r_1)f-(d-r_1)e_1)&\end{aligned} \tag{6.19}$$

and

$$\begin{aligned}\frac{\Gamma_{d-r_1;q;p}(z;\eta/px_1)}{\Gamma_{r_1;q;p}(z;x_1)}\mathcal{S}'_{\eta,x_0,x_1;q;p}(0,ds+d'f-r_1e_1)\frac{\Gamma_{0;q;p}(z;x_1)}{\Gamma_{0;q;p}(z;\eta/px_1)}&\\ =\mathcal{S}_{\eta,x_0\eta/x_1,\eta/px_1;q;p}(0,ds+(d'-r_1)f-(d-r_1)e_1)&\end{aligned} \tag{6.20}$$

Proof. The relevant constraints on the difference operators are manifestly equivalent. □

In particular, we find that the two categories differ merely by a combination of a relabelling of the objects and a gauge transformation, and are in particular isomorphic. This is, of course, hardly surprising; it is simply a noncommutative analogue of the fact that $\mathbb{P}^1\times\mathbb{P}^1$ and F_1 become isomorphic upon blowing up a point.

Lemma 6.3. *We have*

$$T \in \mathcal{S}_{\eta,\eta',x_1;q;p}(0, 2s + 2f - e_1),$$
$$T \in \mathcal{S}'_{\eta,x_0,x_1;q;p}(0, 2s + 3f - e_1).$$

Moreover, for any $v \in \mathbb{Z}s + \mathbb{Z}f + \mathbb{Z}e_1$, we have short exact sequences

$$0 \to T\mathcal{S}_{\eta,\eta',x_1;q;p}(0, v - 2s - 2f + e_1) \to \mathcal{S}_{\eta,\eta',x_1;q;p}(0, v) \to [T^0]\mathcal{S}_{\eta,\eta',x_1;q;p}(0, v) \to 0 \tag{6.21}$$

and

$$0 \to T\mathcal{S}'_{\eta,x_0,x_1;q;p}(0, v - 2s - 3f + e_1) \to \mathcal{S}'_{\eta,x_0,x_1;q;p}(0, v) \to [T^0]\mathcal{S}'_{\eta,x_0,x_1;q;p}(0, v) \to 0. \tag{6.22}$$

Proof. Indeed, T is invariant under the gauge transformation, so is indeed in the given Hom spaces. Similarly, if $[T^0]\mathcal{D} = 0$, then $T^{-1}\mathcal{D}$ lies in the appropriate Hom space as required. □

The reason for the negative signs above on the coefficients of e_1 is that when those coefficients are positive, there is no constraint.

Lemma 6.4. *For any $d, d' \in \mathbb{Z}$, $r_1 \geq 0$,*

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f + r_1e_1) = \mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f)$$
$$\mathcal{S}'_{\eta,x_0,x_1;q;p}(0, ds + d'f + r_1e_1) = \mathcal{S}'_{\eta,x_0;q;p}(0, ds + d'f)$$

Proof. Since we have one inclusion by definition, it will suffice to prove the claim when $d = 0$, $d' = 1$ and when $d = 1$, when it follows from a straightforward calculation. □

Conjugating by the standard elementary transformation gives the following.

Lemma 6.5. *For any $d, d' \in \mathbb{Z}$, $r_1 \geq d$,*

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - r_1e_1) = \theta_p(q^d z/x_1, q\eta/x_1 z; q)_{r_1-d}\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - de_1)$$
$$\mathcal{S}'_{\eta,x_0,x_1;q;p}(0, ds + d'f - r_1e_1) = \theta_p(q^d z/x_1, q\eta/x_1 z; q)_{r_1-d}\mathcal{S}'_{\eta,x_0,x_1;q;p}(0, ds + d'f - de_1).$$

Now, for $d \geq 0$, $r \in \mathbb{Z}$, there is a natural sheaf $V_{d,r}$ on $\mathbb{P}^1$ such that

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - re_1) \subset \Gamma(V_{d,r}(d')), \tag{6.23}$$

with equality for sufficiently large d'. Indeed, this is simply the sheaf associated to the graded module $\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + *f - re_1)$ over a $\mathbb{Z}$-algebra isomorphic to $\mathbb{P}^1$; this has no torsion elements, since a torsion element would be a zero divisor. Similarly, the sheaves $V_{d,r}$ are all torsion-free, and coherent (they are either equal to one or bounded between two sheaves we already understand). We conclude in particular that $V_{d,r}$ is a vector bundle for all $d \geq 0$, $r \in \mathbb{Z}$, of rank $d + 1$.

Lemma 6.6. *The vector bundle $V_{d,r}$ has Hilbert polynomial*

$$h(t; V_{d,r}) = (d + 1)(t + 1) - \begin{cases} 0 & r \leq 0 \\ r(r + 1)/2 & 0 \leq r \leq d \\ (2r - d)(d + 1)/2 & d \leq r \end{cases} \tag{6.24}$$

Moreover, for $0 < r < d$, there is a natural short exact sequence

$$0 \to V_{d-2,r-1}(-2) \to V_{d,r} \to \rho_*\mathcal{L} \to 0, \tag{6.25}$$

where $\mathcal{L}$ is a line bundle on C of degree $2d - r$.

Proof. If $r \leq 0$ or $r \geq d$, this follows from the identification of $V_{d,r}$ with a vector bundle coming from $\mathcal{S}_{\eta,\eta';q;p}$ or $\mathcal{S}'_{\eta,\eta/px_1}$.

Now, multiplication by T gives an injective morphism of graded modules, so that we may identify the image of multiplication by T in $V_{d,r}$ with $V_{d-2,r-1}(-2)$. We may identify the quotient with the image of $[T^0]$, and an easy induction shows that for $d' \geq d$, the image of $[T^0]$ contains all sections of some line bundle of degree $2d + 2d' - r$. (Here we are assuming, as we may, that $d > r > 0$.) We thus find that

$$h(t; V_{d,r}) \geq h(t-2; V_{d-2,r-1}) + 2d + 2d' - r, \tag{6.26}$$

implying by induction the lower bound

$$h(t; V_{d,r}) \geq (d+1)(t+1) - r(r+1)/2. \tag{6.27}$$

(Note that each step in the induction decreases both r and $d - r$ by 1, and thus we will reach the base cases $d = r$ or $r = 0$ before forcing d negative.)

It remains to show that the lower bound is tight. We note that while the the sequence $V_{d,r}$ of vector bundles is decreasing, it cannot decrease *too* quickly, since

$$\begin{aligned}\mathcal{S}_{\eta,\eta',x_1;q;p}(ds + (d'-1)f - (r-1)e_1, ds + d'f - re_1)&\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + (d'-1)f - (r-1)e_1)\\ &\subset \mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - re_1)\\ &\subset \mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - (r-1)e_1)\end{aligned} \tag{6.28}$$

and thus the vector bundles $V_{d,r}$ satisfy

$$V_{d,r-1}(-1) \subset V_{d,r} \subset V_{d,r-1}. \tag{6.29}$$

In particular if we write

$$V_{d,r} = \bigoplus_{0 \leq i \leq d} \mathcal{O}_{\mathbb{P}^1}(e_{d,r,i}) \tag{6.30}$$

with $e_{d,r,0} \geq e_{d,r,1} \geq \cdots \geq e_{d,r,d}$, then

$$e_{d,r-1,i} - 1 \leq e_{d,r,i} \leq e_{d,r-1,i} \tag{6.31}$$

for all $d \geq i \geq 0$, $r \in \mathbb{Z}$.

Now, for $r \geq d$, the identification of $V_{d,r}$ with a family of Hom spaces of $\mathcal{S}'$, together with the known basis of these spaces lets us compute the corresponding degree sequence. We thus find that $e_{d,r,i} = d - r - i$ for $0 \leq i \leq d \leq r$. For $r \leq 0$, the situation is somewhat more subtle, since the degree sequence depends on the parameters in general. Luckily, we will only need to understand the differences $e_{d,0,i} - e_{d,d,i}$. Although these still depend on the parameters (since the sequence $e_{d,d,i}$ doesn't), it turns out that the dependence on the parameters is relatively mild. In particular, although the sequence $e_{d,0,i} - e_{d,d,i}$ depends on the parameters, the *multiset* doesn't! Indeed, an easy induction using Lemma 4.22 tells us that the sequence $e_{d,0,i} - e_{d,d,i} = e_{d,0,i} + i$ is a permutation of the sequence $0, 1, \ldots, d$ corresponding to the generic case.

The significance of this fact is that as we increase r from 0 to d, each degree can only decrease by 1. In particular, one of the degrees is decreased at every step, one at all but one steps, etc. In particular, the first r steps must decrease the largest difference every time, the second largest difference at least $r - 1$ times, etc. But this implies

$$\sum_{0 \leq i \leq d} e_{d,r,i} \leq \sum_{0 \leq i \leq d} e_{d,0,i} - r(r-1)/2, \tag{6.32}$$

giving the other required inequality on $\deg(V_{d,r})$. □

Remark. This corresponds to the following interpretation of $\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds+d'f-r_1e_1)$ for generic parameters: it is the subspace of $\mathcal{S}_{\eta,\eta';q;p}(0, ds+d'f)$ consisting of operators for which $[T^r]\mathcal{D}$ vanishes at $q^{-k}x_1$ for $r \le k < r_1$. (These are precisely the points where the relevant gauge transformation introduces poles, and the total number of conditions is correct.) In the same way, $\mathcal{S}'_{\eta,x_0,x_1;q;p}(0, ds+d'f-r_1e_1)$ is generically the subspace of $\mathcal{S}'_{\eta,x_0;q;p}(0, ds+d'f)$ consisting of operators for which $[T^r]\mathcal{D}$ vanishes at $q^{-k}x_1$ for $r \le k < r_1$.

Corollary 6.7. *If $d' \ge d \ge r_1 \ge 0$, then*

$$\dim(\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds+d'f-r_1e_1)) = (d+1)(d'+1) - r_1(r_1+1)/2. \tag{6.33}$$

Proof. Since $\mathcal{S}_{\eta,\eta';q;p}(0, ds+d'f) = \hat{\mathcal{S}}_{\eta,\eta';q;p}(0, ds+d'f)$ for $d' \ge d-1$, the same applies when $r \ge 0$. But then we may identify the given space with the space of global sections of a bundle V_{d,r_1}, and we need simply show that this bundle is acyclic. This follows by an easy induction from the leading coefficient short exact sequence, with base case $r_1 = 0$ or $r_1 = d$. □

Lemma 6.8. *Suppose $\mathcal{D} \in \mathcal{S}_{\eta,\eta';q;p}(0, ds+d'f)$ is such that*

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(ds+d'f, (d+1)s+(d'+1)f)\mathcal{D} \subset \mathcal{S}_{\eta,\eta',x_1;q;p}(0, (d+1)s+(d'+1)f-r_1e_1). \tag{6.34}$$

Then $\mathcal{D} \in \mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds+d'f-r_1e_1)$.

Proof. By assumption, for any operator

$$\mathcal{D}_{q^{-d}\eta;q;p}(b(z)) \in \mathcal{S}_{\eta,\eta',x_1;q;p}(ds+d'f, (d+1)s+(d'+1)f), \tag{6.35}$$

we have

$$\mathcal{D}_{q^{-d}\eta;q;p}(b(z))\mathcal{D} \in \mathcal{S}_{\eta,\eta',x_1;q;p}(0, (d+1)s+(d'+1)f-r_1e_1). \tag{6.36}$$

Thus

$$\frac{\Gamma_{d+1-r_1;q;p}(z;\eta/px_1)}{\Gamma_{r_1}(x;z_1)}\mathcal{D}_{q^{-d}\eta;q;p}(b(z))\mathcal{D}\frac{\Gamma_0(x;z_1)}{\Gamma_{0;q;p}(z;\eta/px_1)} \in \mathcal{S}'_{\eta,\eta'/px_1}(0, (d+1)s+(d+d'-r_1+2)f), \tag{6.37}$$

and it remains to show

$$\mathcal{D}' := \frac{\Gamma_{d-r_1;q;p}(z;\eta/px_1)}{\Gamma_{r_1}(x;z_1)}\mathcal{D}\frac{\Gamma_0(x;z_1)}{\Gamma_{0;q;p}(z;\eta/px_1)} \in \mathcal{S}'_{\eta,\eta'/px_1}(0, ds+(d+d'-r_1)f). \tag{6.38}$$

Since

$$\frac{\Gamma_{d+1-r_1;q;p}(z;\eta/px_1)}{\Gamma_{r_1}(x;z_1)}\mathcal{D}_{q^{-d}\eta;q;p}(b(z))\frac{\Gamma_{r_1}(x;z_1)}{\Gamma_{d-r_1;q;p}(z;\eta/px_1)} = \mathcal{D}_{q^{-d}\eta;q;p}(\theta_p(q^{r_1-d}\eta/x_1z)b(z)), \tag{6.39}$$

we reduce to showing that

$$\mathcal{D}_{q^{-d}\eta;q;p}(\theta_p(q^{r_1-d}\eta/x_1z)b(z))\mathcal{D}' \in \mathcal{S}'_{\eta,\eta'/px_1}(0, (d+1)s+(d+d'-r_1+2)f) \tag{6.40}$$

for all b implies $\mathcal{D}' \in \mathcal{S}'_{\eta,\eta'/px_1}(0, ds+(d+d'-r_1)f)$. This follows immediately from the adjoint of Lemma 5.8 above: simply choose b so that

$$b_2(z) = \theta_p(q^{r_1-d}\eta/x_1z)b(z) \tag{6.41}$$

satisfies the hypotheses of that Lemma. □

Theorem 6.9. *The Fourier transform*

$$\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f) \cong \mathcal{S}_{\eta',\eta;q;p}(0, d's + df) \tag{6.42}$$

restricts to isomorphisms

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - r_1e_1) \cong \mathcal{S}_{\eta',\eta,x_1;q;p}(0, d's + df - r_1e_1) \tag{6.43}$$

for all integers r_1.

Proof. Note first that Lemma 6.8 implies that if the result holds for $ds + d'f - r_1e_1$, then it holds for $(d-1)s + (d'-1)f - r_1e_1$. We may thus reduce to the case $\min(d, d') \geq r_1$. If $r_1 \leq 0$, there is nothing to prove. Otherwise, we claim that

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds+d'f-r_1e_1) = \mathcal{S}_{\eta,\eta',x_1;q;p}(s+f-e_1, ds+d'f-r_1e_1)\mathcal{S}_{\eta,\eta',x_1;q;p}(0, s+f-e_1). \tag{6.44}$$

Indeed, the right-hand side accounts for all possible leading coefficients, while

$$T \in \mathcal{S}_{\eta,\eta',x_1;q;p}(s+f-e_1, 2s+2f-e_1)\mathcal{S}_{\eta,\eta',x_1;q;p}(0, s+f-e_1), \tag{6.45}$$

by the adjoint of Lemma 3.3. We thus reduce to showing the claim in the case $d = d' = 1$.

Now,

$$\mathcal{D}_{\eta;q;p}(b) \in \mathcal{S}_{\eta,\eta',x_1;q;p}(0, s+f-e_1) \tag{6.46}$$

iff

$$\Gamma_{p,q}(qz/x_1, q\eta/x_1z)^{-1}\mathcal{D}_{\eta;q;p}(b)\Gamma_{p,q}(z/x_1, q\eta/x_1z) = \mathcal{D}_{\eta;q;p}(b/\theta_p(z/x_1)) \in \mathcal{S}'_{\eta,\eta'/px_1}(0, s+f). \tag{6.47}$$

In other words, $\mathcal{D}_{\eta;q;p}(b) \in \mathcal{S}_{\eta,\eta',x_1;q;p}(0, s+f-e_1)$ if $b(x_1) = 0$. This condition is independent of η, η', and is thus preserved by the Fourier transform. □

Corollary 6.10. *The family* $\mathcal{S}_{\eta,\eta',x_1;q;p}$ *of* $\mathbb{Z}^3$*-algebras is flat.*

Proof. We need to show that the dimension of

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - r_1e_1) \tag{6.48}$$

depends only on d, d', r_1. The Theorem tells us that

$$\dim \mathcal{S}_{\eta,\eta',x_1;q;p}(0, ds + d'f - r_1e_1) = \dim \mathcal{S}_{\eta',\eta,x_1;q;p}(0, d's + df - r_1e_1), \tag{6.49}$$

and we may thus reduce to the case $d' \geq d$. If $r_1 \geq d$, then the space is essentially a Hom space in $\mathcal{S}'_{\eta,\eta'/px_1;q;p}$, and is thus flat. Similarly, if $r_1 \leq 0$, then the space is just $\mathcal{S}_{\eta,\eta';q;p}(0, ds + d'f)$, so again is flat. Finally, if $0 \leq r_1 \leq d$, then $\mathcal{S}_{\eta,\eta';q;p}(0, ds+d'f-r_1e_1)$ is the space of global sections of an acyclic vector bundle of Euler characteristic $(d+1)(d'+1) - r_1(r_1+1)/2$, and the claim follows immediately. □

Applying the standard elementary transformation gives the following.

Corollary 6.11. *The family* $\mathcal{S}'_{\eta,x_0,x_1;q;p}$ *of* $\mathbb{Z}^3$*-algebras is flat, and there is an isomorphism*

$$\mathcal{S}'_{\eta,x_0,x_1;q;p} \cong \mathcal{S}'_{x_0\eta/x_1,x_1,x_0;q;p}, \tag{6.50}$$

which on objects acts as $ds + d'f - r_1e_1 \mapsto (d'-r_1)s + d'f - (d'-d)e_1$.

This of course is just a noncommutative analogue of the fact that if we blow up $\mathbb{P}^2$ in two distinct points, the order of the two points is irrelevant. (Of course, it is somewhat stronger than that even in the commutative case, as we have made no assumption on x_0 and x_1.)

The adjoint symmetry also extends to this case. The key observation is that we should take

$$\Gamma_{r;1/q;p}(z;x) = \Gamma_{1-r;q;p}(z;x)^{-1}, \tag{6.51}$$

insofar as both sides have the same behavior under $T = T_{1/q}^{-1}$ and $r \mapsto r+1$. The following is then an immediate consequence of the adjoint symmetries for $\mathcal{S}_{\eta,\eta';q;p}$ and $\mathcal{S}'_{\eta,x_0;q;p}$.

Proposition 6.12. *The adjoint isomorphism*

$$\mathcal{S}_{\eta,\eta';q;p} \cong \mathcal{S}^{\mathrm{op}}_{\eta,\eta';1/q;p} \tag{6.52}$$

extends to an isomorphism

$$\mathcal{S}_{\eta,\eta',x_1;q;p} \cong \mathcal{S}^{\mathrm{op}}_{\eta,\eta',x_1;1/q;p} \tag{6.53}$$

which on objects acts as $ds + d'f - r_1e_1 \mapsto (2s+2f-e_1) - (ds+d'f-r_1e_1)$.

Remark. Similarly, for $\mathcal{S}'_{\eta,x_0,x_1;q;p}$, the adjoint acts as $ds+d'f-r_1e_1 \mapsto (2s+3f-e_1)-(ds+d'f-r_1e_1)$.

In order to extend this construction to algebraic curves, it will be helpful to pin down a relatively small set of generators.

Proposition 6.13. *The category* $\mathcal{S}_{\eta,\eta',x_1;q;p}$ *is generated by elements of degrees* e_1, $f-e_1$, $s-e_1$, f, s, $s+f-e_1$.

Proof. Consider a Hom space of degree $ds+d'f-r_1e_1$. If $r_1 < 0$, then we may factor out morphisms of degree e_1 to reduce to $r_1 = 0$, for which the claim is immediate. Similarly, if $r_1 > d$, we may factor out morphisms of degree $f-e_1$, while if $d' > d$, we may apply the Fourier transform. We thus see that we need only consider Hom spaces of degree $ds+d'f-r_1e_1$ with $d' \geq d \geq r_1 \geq 0$.

Now, the degrees we have reduced to considering form a semigroup, and it is easy to see the semigroup is generated by f, $s+f-e_1$, and $s+f$. The first two degrees are already accounted for by our generators, while for the third, we may write

$$\begin{aligned}\mathcal{S}_{\eta,\eta',x_1;q;p}(0,s+f) \supset\ & \mathcal{S}_{\eta,\eta',x_1;q;p}(f,s+f)\mathcal{S}_{\eta,\eta',x_1;q;p}(0,f) + \mathcal{S}_{\eta,\eta',x_1;q;p}(s,s+f)\mathcal{S}_{\eta,\eta',x_1;q;p}(0,s) \\ & + \mathcal{S}_{\eta,\eta',x_1;q;p}(s+f-e_1,s+f)\mathcal{S}_{\eta,\eta',x_1;q;p}(0,s+f-e_1)\end{aligned} \tag{6.54}$$

Now, we have already seen that the first two terms span $\mathcal{S}_{\eta,\eta',x_1;q;p}(0,s+f) = \mathcal{S}_{\eta,\eta';q;p}(0,s+f)$ unless $q=1$, $\eta=\eta'$, when their span can be identified with the space of $BC_1(\eta)$-symmetric theta functions. Since the third term then consists of those theta functions with the same multiplier that vanish at x_1, it follows that $\mathcal{S}_{\eta,\eta',x_1;q;p}(0,s+f)$ is indeed contained in the generated subcategory.

We next observe that $\mathcal{S}_{\eta,\eta',x_1;q;p}(0,2s+2f-e_1)$ is contained in the generated subcategory. We first note that by Lemma 2.3, for every operator in $\mathcal{S}_{\eta,\eta',x_1;q;p}(0,2s+2f-e_1)$, there is an operator in $\mathcal{S}_{\eta,\eta',x_1;q;p}(0,s+f-e_1)\mathcal{S}_{\eta,\eta',x_1;q;p}(s+f)$ with the same leading coefficient. Since it follows immediately from Lemma 3.3 that

$$T \in \mathcal{S}_{\eta,\eta',x_1;q;p}(0,s+f-e_1)\mathcal{S}'_{\eta,\eta',x_1;q;p}(s+f), \tag{6.55}$$

the claim follows.

Now, consider an element $ds+d'f-r_1e_1$ of the semigroup. If $r_1 = 0$, then it follows from the definition of $\mathcal{S}_{\eta,\eta';q;p}$ that every operator of this degree is in the generated subcategory. Similarly,

if $r_1 = d$, then the claim follows by considering $\mathcal{S}'_{\eta,\eta'/px_1;q;p}$. Finally, if $d > r_1 > 0$, then we may subtract $2s + 2f - e_1$ without leaving the semigroup. We thus find that every element of $\mathcal{S}_{\eta,\eta',x_1;q;p}(0,(d-2)s+(d'-2)f-(r_1-1)e_1)$ is contained in the subcategory, so that it suffices to show that

$$\begin{aligned}&\mathcal{S}_{\eta,\eta',x_1;q;p}((d-2)s+(d'-2)f-(r_1-1)e_1, ds+d'f-r_1e_1)\\ &\circ \mathcal{S}_{\eta,\eta',x_1;q;p}(0,(d-2)s+(d'-2)f-(r_1-1)e_1) = \mathcal{S}_{\eta,\eta',x_1;q;p}(0,ds+d'f-r_1e_1). \end{aligned} \tag{6.56}$$

This certainly accounts for every element with leading coefficient 0, and by Lemma 2.3 accounts for every possible leading coefficient. □

Remark. If $q \neq 1$ or $\eta \neq \eta'$, then the generators of degree $s + f - e_1$ are redundant. If $q \neq 1$, then the generators of degree s and f are also redundant.

It immediately follows that for $\mathcal{S}'_{\eta,x_0,x_1;q;p}$, we can generate the category by elements of degree s, e_1, $f - e_1$, f, $s + f - e_1$, and $s + f$. Thus to define $\mathcal{S}'_{\eta,x_0,x_1;q;C}$, it will suffice to determine the operators corresponding to morphisms of degree e_1, $f - e_1$ and $s + f - e_1$. Now, the operators of degree e_1 are just scalar multiples of 1 (since they must be contained in the space of operators of degree 0), while for $f - e_1$ and $s + f - e_1$, the only constraint is that the leading coefficient must vanish at x_1 (in addition to the already forced divisor). Of course, we can then also define $\mathcal{S}_{\eta,\eta',x_1;q;C} := \mathcal{S}'_{\eta,\eta'/x_1,\eta/x_1;q;C}$, and can then obtain an elliptic version of $\mathcal{S}_{\eta,\eta';q;p}$ by taking the appropriate $\mathbb{Z}^2$-subalgebra. The result here depends on the choice of x_1, but changing x_1 gives an isomorphic category.

When $q = 1$, $\eta \neq \eta'$, the generators of degree $s + f - e_1$ are redundant, and we moreover see that (assuming we choose the other generators in a way independent of the source object) the only relations are that the generators commute. This gives us the following, not unexpected, result.

Proposition 6.14. *Let C be a genus 1 curve (over an algebraically closed field) embedded in $\mathbb{P}^1 \times \mathbb{P}^1$ by a pair of (nonisomorphic) line bundles η, η' of degree 2, and let X be the surface obtained from $\mathbb{P}^1 \times \mathbb{P}^1$ by blowing up the point $x_1 \in C$. The subcategory of $\mathrm{Coh}(X)$ generated by line bundles is equivalent to $\mathcal{S}_{\eta,\eta',x_1;1;C}$, in such a way that the restriction to $C \subset X$ corresponds to taking leading coefficients.*

Proof. For each $d, d', r_1 \in \mathbb{Z}$, let $\mathcal{D}_{ds+d'f-r_1e_1}$ be the divisor on C given by

$$\mathcal{D}_{ds+d'f-r_1e_1} = [\eta'/x_1]^{d'-r_1}[\eta/x_1]^{r_1-d}[x_1\eta/\eta']^{d'-d-r_1}R^d_\eta. \tag{6.57}$$

This divisor plays two roles. First, the leading coefficient of an element of $\mathcal{S}_{\eta,\eta',x_1;1;C}(0, ds + d'f - r_1e_1)$ is naturally a global section of the line bundle $\mathcal{L}(\mathcal{D}_{ds+d'f-r_1e_1})$. (This is really a statement about $\mathcal{S}'_{\eta,\eta'/x_1,\eta/x_1;1;C}$, where it follows readily from the corresponding statement on F_1.) Second, if $\mathcal{L}_{ds+d'f-r_1e_1}$ is a line bundle of the specified divisor class on X, then

$$\mathcal{L}_{ds+d'f-r_1e_1}|_C \cong \mathcal{L}(\mathcal{D}_{ds+d'f-r_1e_1}). \tag{6.58}$$

(Indeed, $\mathcal{D}_{ds+d'f-r_1e_1}$ is a representative of the divisor class $\eta^{d'}\eta'^{d}x_1^{-r_1}$.)

Now, choose a representative of each divisor class on X, and for each such representative choose an identification

$$\mathcal{L}_{ds+d'f-r_1e_1}|_C \cong \mathcal{L}(\mathcal{D}_{ds+d'f-r_1e_1}). \tag{6.59}$$

Then there is a functor from $\mathcal{S}_{\eta,\eta',x_1;1;C}$ to this subcategory (acting in the obvious way on objects) making it naturally an equivalence. Indeed, this identification removes any freedom in choosing

how generators map, and since the only relations in $\mathcal{S}_{\eta,\eta',x_1;1;C}$ are that the generators commute, this gives a functor. It is moreover straightforward to verify that the generators are algebraically independent as functions on X, and that the Hom spaces on both sides have the same dimension. □

Remark. When $\eta = \eta'$, we find in a similar way that $\mathcal{S}_{\eta,\eta,x_1;1;C}$ can be identified with a non-full subcategory of $\mathrm{Coh}(X)$, where now X is obtained by blowing up the point x_1 of C, embedded in F_2 via η.

Another use for this small set of generators is that it allows us to define morphisms from $\mathcal{S}_{\eta,\eta',x_1;q;p}$. Of course, we have already dealt with most of the morphisms we will need, with the one caveat that some of them are only truly defined up to a scalar gauge (as the Fourier transform and the adjoint both require us to a choose a certain system of elliptic functions). There is one more that is worth considering, this time no longer an equivalence.

Proposition 6.15. *Suppose q is a primitive r-th root of unity. There is a functor*

$$\mathcal{S}_{\eta^r,\eta'^r,x_1^r;1;p^r} \to \mathcal{S}_{\eta,\eta',x_1;q;p} \tag{6.60}$$

which on objects acts as multiplication by r, and on morphisms acts as

$$\sum_{0\le k\le d} c_k(z)T^k \mapsto \sum_{0\le k\le d} c_k(z^r)T^{rk}. \tag{6.61}$$

Proof. This operation clearly respects composition, so it remains only to show that it takes generators to morphisms. For degree e_1, this is immediate, as the functor takes the operator 1 to the operator 1. For degree $f - e_1$, we note that any element of

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0, rf - re_1) \tag{6.62}$$

is a composition of elements of degree $f - e_1$, and is thus proportional to

$$\prod_{0\le i<r} \theta_p(z/q^r x_1, \eta/q^r x_1 z) = \theta_{p^r}(z^r/x_1^r, \eta^r/x_1^r z^r), \tag{6.63}$$

from which the claim follows. Now, this implies in particular that

$$\theta_{p^r}(z^r/x_1^r, \eta^r/x_1^r z^r) \in \mathcal{S}_{\eta,\eta';q;p}(0, rf); \tag{6.64}$$

it immediately follows that

$$\theta_{p^r}(z^r/x^r, \eta^r/x^r z^r) \in \mathcal{S}_{\eta,\eta';q;p}(0, rf) \tag{6.65}$$

for any x, so that the putative functor acts correctly on elements of degree f.

Now, for the generator of degree $s - e_1$, we first apply the elementary transform and consider

$$\mathcal{S}'_{\eta,\eta'/px_1,\eta/px_1;q;p}(0, rs) = \mathcal{S}'_{\eta,\eta'/px_1;q;p}(0, rs) = \mathcal{S}'_{\eta,\eta'/px_1,\eta'/px_1;q;p}(0, rs). \tag{6.66}$$

Thus up to a gauge transformation, we are considering elements of

$$\mathcal{S}_{\eta,\eta,x_1\eta/\eta';q;p}(0, rs - re_1) \subset \mathcal{S}_{\eta,\eta;q;p}(0, rs) \tag{6.67}$$

By our known results on saturation in $\mathcal{S}_{\eta,\eta;q;p}$, any such operator factors as

$$\mathcal{S}_{\eta,\eta;q;p}(rs - rf, rs)\mathcal{S}_{\eta,\eta;q;p}(0, rs - rf) \tag{6.68}$$

and thus is a left multiple of the operator

$$\frac{z^r}{\theta_{p^r}(z^{2r}/\eta^r)}(1-T^r). \tag{6.69}$$

It follows (without needing to consider details of the gauge transformations) that any element of

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0,rs-re_1) \tag{6.70}$$

has the form $c_0(z)+c_r(z)T^r$ for suitable functions $c_0(z)$, $c_r(z)$. Any such operator is also a composition of operators of degree $s-e_1$, and it is straightforward to compute the resulting extreme coefficients. (Indeed, via the Fourier transform, we see that this is the same as the calculation for degree $rf-re_1$.)

We thus find that the (unique up to scalars) global section of

$$\mathcal{S}_{\eta,\eta',x_1;q;p}(0,rs-re_1) \tag{6.71}$$

has the form

$$\frac{z^r}{\theta_{p^r}(z^{2r}/\eta^r)}\theta_p(z^r/x_1^r,\eta'^r/z^r x_1^r)-\frac{z^r}{\theta_{p^r}(z^{2r}/\eta^r)}\theta_p(\eta^r/z^r x_1^r,z^r\eta'^r/\eta^r x_1^r)T^r \tag{6.72}$$

as required. Again, we find that the same operator gives a global section of $\mathcal{S}_{\eta,\eta';q;p}(0,rs)$ even if we change x_1, and thus obtain a full set of generators of degree s.

Since the generator of degree $s+f-e_1$ is redundant unless $\eta^r=\eta'^r$, flatness implies that such elements map to morphisms in general. □

Remark. For $\mathcal{S}_{\eta,\eta',x_1;q;C}$, the corresponding fact is that if $q\in\mathrm{Pic}^0(C)$ has order r and $\phi:C\to C'$ is the corresponding quotient morphism of degree r, then there is a functor from $\mathcal{S}_{\phi_*\eta,\phi_*\eta',\phi_*x_1;1;C'}$ to $\mathcal{S}_{\eta,\eta',x_1;q;C}$. We will see below that this essentially makes $\mathcal{S}_{\eta,\eta',x_1;q;C}$ a (degenerate) Azumaya algebra over the commutative surface corresponding to $\mathcal{S}_{\phi_*\eta,\phi_*\eta',\phi_*x_1;1;C'}$.

7 General blowups

The main difficulty in extending the above construction to more general blowups is that the above condition for the $K^2=7$ case is very difficult to extend in a uniform (and flat) manner. For the commutative case, the difficulty arises when blowing up the same point multiple times, as then we need to consider functions vanishing with multiplicity along jets. In the noncommutative case, this becomes even worse, as any two points in the same $q^{\mathbb{Z}}$ orbit will encounter similar issues, and would seem to require us to define the conditions inductively.

If the points being blown up are in sufficiently general position, then there is no difficulty: we expect the conditions at the different points to be entirely independent in such a case. The idea, then, is to deal with the general case by taking a suitably defined limit from the generic case. Of course, in order to make sense of such a limit, we will need to ensure that our $\mathbb{Z}^{m+2}$-algebras have a suitable sheaf structure.

In general, let F be a flat sheaf on an integral scheme S, and suppose we are given a finite-dimensional subspace V of the generic fiber of F. (For instance, F might be a Hom space in $\mathcal{S}$ or $\mathcal{S}'$, and V the subspace corresponding to a Hom space in the generic iterated blowup.) Then there is a natural extension of V to a subsheaf F_V of F having that generic fiber. Indeed, any such sheaf will have $\Gamma(U,F_V)\subset\Gamma(U,F)$ for any open subset U, and the restriction of any element of $\Gamma(U,F_V)$ to

the generic fiber will be an element of V. Since the presheaf of sections satisfying those conditions is a sheaf, this gives the desired natural extension (which is clearly maximal among all subsheaves with generic fiber V). We note that F/F_V is torsion-free, since otherwise we could replace F_V by the preimage of the torsion subsheaf of F/F_V without affecting the generic fiber. (Conversely, any subsheaf with torsion-free quotient is the natural extension of its generic fiber.) Also, if F is an algebra, then the product structure induces a multiplication on canonical extensions; if V and W are two subspaces of the generic fiber, there is a natural multiplication $F_V \otimes F_W \to F_{VW}$.

Naturally, we will want to prove that the canonical extension in our case is flat (as the whole point of the exercise is to obtain a flat family). In fact, we want something stronger. Our objective, after all, is not just to construct a flat family of $\mathbb{Z}^{m+2}$-algebras, but also faithful representations of the fibers of that family in terms of difference operators. The difficulty is that, in terms of the above setup, even if F_V is flat, the fibers of F_V may fail to inject in the fibers of F. The kernel of the map on fibers comes from Tor_1 of the quotient F/F_V, and thus we see that we also want to insist that F/F_V be flat. (In fact, it is then redundant to require F_V to be flat.) We will say that $V \subset F_{k(S)}$ "has coflat extension in F" if F/F_V is flat, and similarly say that the extension F_V is "coflat".

We note here that if $V \subset W \subset F_{k(S)}$, and W has coflat extension in F, then V has coflat extension in F iff it has coflat extension in F_W. We will thus mainly omit mention of the ambient sheaf, though we caution the reader that varying the ambient sheaf in any other way can have a significant effect on the extension.

For any pair (C, q) with C a smooth genus 1 curve and $q \in \mathrm{Pic}^0(C)$, let $\mathcal{E}ll\mathcal{D}iff_{q;C}$ denote the algebra of elliptic difference operators; that is $\mathcal{E}ll\mathcal{D}iff_{q;C}$ is the algebra of finite linear combinations $\sum_{0\le k} c_k(z)T^k$, where $c_k(z)$ are elements of the function field of C, and conjugation by T acts via translation by q. These algebras fit in a natural way into a flat family of algebras over the moduli stack of pairs (C, q). This lets us define for any $m \ge 0$ a flat family of $\mathbb{Z}^{m+2}$-algebras over the moduli stack of tuples $(C, q, \eta, x_0, x_1, \ldots, x_m)$ with C smooth of genus 1, $q \in \mathrm{Pic}^0(C)$, $\eta \in \mathrm{Pic}^2(C)$, and $x_i \in C$, in which every Hom space is $\mathcal{E}ll\mathcal{D}iff_{q;C}$. (This moduli stack is not a scheme, but is an integral algebraic space, as can be seen by adjoining suitable level structures; as a result, the canonical extension construction still applies.) It will be convenient on occasion to encode the parameters $\eta, x_0, \ldots, x_m$ in a single object, namely the homomorphism

$$\rho : \langle s, f, e_1, \ldots, e_m\rangle \to \mathrm{Pic}(C) \tag{7.1}$$

given by

$$\rho(s) = x_0, \quad \rho(f) = \eta, \quad \rho(e_i) = x_i. \tag{7.2}$$

Note that in the commutative situation, ρ is just the natural restriction map $\mathrm{Pic}(X) \to \mathrm{Pic}(C)$.

We then define a family of $\mathbb{Z}^{m+2}$-algebras $\mathcal{S}'_{\rho;q;C} = \mathcal{S}'_{\eta,x_0,\ldots,x_m;q;C}$ as the canonical extension in the above $\mathbb{Z}^{m+2}$-algebra version of $\mathcal{E}ll\mathcal{D}iff_{q;C}$ of the following generic category:

$$\begin{aligned}&\mathcal{S}'_{\eta,x_0,x_1,\ldots,x_m;q;C}(d_1s + d'_1f - r_{11}e_1 - \cdots - r_{1m}e_m, d_2s + d'_2f - r_{21}e_1 - \cdots - r_{2m}e_m)\\ &= \bigcap_{1\le i\le m} \mathcal{S}'_{\eta,x_0,x_i;q;C}(d_1s + d'_1f - r_{1i}e_i, d_2s + d'_2f - r_{2i}e_i).\end{aligned} \tag{7.3}$$

In other words, given an open subset U of the moduli stack, the elements of

$$\Gamma(U, \mathcal{S}'_{\eta,x_0,x_1,\ldots,x_m;q;C}(d_1s + d'_1f - r_{11}e_1 - \cdots - r_{1m}e_m, d_2s + d'_2f - r_{21}e_1 - \cdots - r_{2m}e_m)) \tag{7.4}$$

are precisely the holomorphic families of difference operators parametrized by U which are generically in the above intersection. In light of Lemma 6.2, we also define a family $\mathcal{S}_{\eta,\eta',x_1,\ldots,x_m;q;C}$

by

$$\begin{aligned}&\mathcal{S}_{\eta,\eta',x_1,\dots,x_m;q;C}(d_1s+d_1'f-r_{11}e_1-\cdots-r_{1m}e_m, d_2s+d_2'f-r_{21}e_1-\cdots-r_{2m}e_m)\\ :=\ &\mathcal{S}'_{\eta,\eta'/x_1,\eta/x_1,x_2,\dots,x_m;q;C}(d_1s+(d_1+d_1'-r_{11})f-(d_1-r_{11})e_1-r_{12}e_2\cdots-r_{1m}e_m,\\ &\qquad d_2s+(d_2+d_2'-r_{21})f-(d_2-r_{11})e_1-r_{22}e_2\cdots-r_{2m}e_m).\end{aligned}\tag{7.5}$$

Again, it will occasionally be useful to encode the parameters other than q and C in a single homomorphism

$$\rho:\langle s,f,e_1,\dots,e_m\rangle\to \mathrm{Pic}(C),\tag{7.6}$$

with

$$\rho(s)=\eta',\quad \rho(f)=\eta,\quad \rho(e_i)=x_i.\tag{7.7}$$

Since the canonical extension only depends on the ambient sheaf and the generic fiber, the following is immediate.

Proposition 7.1. *For any permutation $\pi\in S_m$, there is an isomorphism (of sheaves of $\mathbb{Z}^{m+2}$-algebras)*

$$\mathcal{S}'_{\eta,x_0,x_1,\dots,x_m;q;C}\cong\mathcal{S}'_{\eta,x_0,x_{\pi(1)},\dots,x_{\pi(m)};q;C}\tag{7.8}$$

acting linearly on objects by $s,f\mapsto s,f$ and $e_i\mapsto e_{\pi(i)}$.

Proof. Indeed, the corresponding generic fibers are manifestly isomorphic. □

Remark. Note that in terms of the homomorphism ρ, the action on parameters is just $\rho\mapsto\rho\circ\pi$, where π acts on $\langle s,f,e_1,\dots,e_m\rangle$ in the same way as it does on objects.

The twisting symmetries also carry over immediately, as the corresponding Hom sheaves are equal as subsheaves of $\mathcal{E}ll\mathcal{D}iff_{q;C}$.

Proposition 7.2. *There are natural isomorphisms*

$$\begin{aligned}\mathcal{S}'_{\eta,x_0,x_1,\dots,x_m;q;C}&\cong\mathcal{S}'_{q^{-d}\eta,q^{d-d'}x_0,q^{-r_1}x_1,\dots,q^{-r_m}x_m;q;C}\\ \mathcal{S}_{\eta,\eta',x_1,\dots,x_m;q;C}&\cong\mathcal{S}_{q^{-d}\eta,q^{-d'}\eta',q^{-r_1}x_1,\dots,q^{-r_m}x_m;q;C}\end{aligned}$$

which act on objects as $v\mapsto v-ds-d'f+r_1e_1+\cdots+r_me_m$.

Remark. The action on parameters has a uniform description in terms of ρ and the intersection pairing: translating by $D_0=ds+d'f-r_1e_1-\cdots-r_me_m$ has the effect of replacing ρ by

$$\rho':D\mapsto q^{-D\cdot D_0}\rho(D).\tag{7.9}$$

More generally, we can construct isomorphisms between canonical extensions by giving isomorphisms between ambient sheaves that restrict to isomorphisms between the respective generic fibers. The simplest is the adjoint.

Proposition 7.3. *There is a natural isomorphism*

$$\mathcal{S}'_{\rho;q;C}\cong(\mathcal{S}'_{\rho;q^{-1};C})^{op}\tag{7.10}$$

which acts on objects as

$$D\mapsto 2s+3f-e_1-\cdots-e_m-D\tag{7.11}$$

and on operators as the formal adjoint in $\mathcal{S}'_{\eta,x_0;q;p}$.

Proof. This follows immediately from the case $m = 1$. □

Remark. There is a technical issue we have glossed over here and below, namely in this case that the adjoint is only determined once we have chosen a countable infinity of elliptic functions with specified divisors, and it is not clear that one can make such choices on an open covering of the parameter space. This is not a serious issue, though, as until we return to considering individual fibers, we only ever consider finitely many objects at a time, and thus there is no difficulty in making the requisite choices Zariski locally. In addition, if we base change to give ourselves a generic point v, then we could simply insist that all the chosen functions evaluate to 1 at v.

Proposition 7.4. *There is an isomorphism of sheaves of $\mathbb{Z}^{m+2}$-algebras*

$$\mathcal{S}'_{\eta,x_0,x_1,\dots,x_m;q;C} \cong \mathcal{S}'_{\eta,x_0\eta/x_1x_2,\eta/x_2,\eta/x_1,x_3,\dots,x_m;q;C} \tag{7.12}$$

acting on objects as

$$\begin{aligned} &ds + d'f - r_1e_1 - r_2e_2 - r_3e_3 - \cdots - r_me_m \\ &\mapsto ds + (d + d' - r_1 - r_2)f - (d - r_2)e_1 - (e - r_1e_2) - r_3e_3 - \cdots - r_me_m. \end{aligned} \tag{7.13}$$

Proof. Since the generic point of the moduli space lives over a characteristic 0 field, we can represent it as a configuration on a complex elliptic curve, and thus to understand the generic fiber of $\mathcal{S}'_{\eta,x_0,x_1,\dots,x_m;q;C}(0, ds + d'f - r_1e_1 - \cdots - r_me_m)$, it suffices to understand

$$\bigcap_{1\le i\le m} \mathcal{S}'_{\eta,x_0,x_i;q;p}(0, ds + d'f - r_ie_i), \tag{7.14}$$

as this differs only by the appropriate gauge transformation. (We can even choose a representative such that $|q| < 1$, so that the gauge transformation is by actual Γ functions.) Now, since the configuration is generic,

$$\mathcal{S}'_{\eta,x_0,x_i;q;p}(0, ds + d'f - r_ie_i) \subset \mathcal{S}'_{\eta,x_0;q;p}(0, ds + d'f) \tag{7.15}$$

is the subspace of operators such that $[T^r]\mathcal{D}$ vanishes at $q^{-k}x_i$, for $r \le k < r_i$. The conditions for $i > 1$ are preserved under the isomorphism

$$\mathcal{S}'_{\eta,x_0,x_1;q;p}(0, ds + d'f - r_1e_1) \cong \mathcal{S}_{\eta,x_0\eta/x_1,\eta/px_1;q;p}(0, ds + (d' - r_1)f - (d - r_1)e_1), \tag{7.16}$$

as this involves gauging by a family of cocycles which are holomorphic and nonzero at all of the points $x_2, \dots, x_m$. It follows that this gauge transformation maps into

$$\mathcal{S}_{\eta,x_0\eta/x_1,\eta/px_1;q;p}(0, ds+(d'-r_1)f-(d-r_1)e_1) \cap \bigcap_{2\le i\le m} \mathcal{S}_{\eta,x_0\eta/x_1,x_i;q;p}(0, ds+(d'-r_1)f-r_ie_i). \tag{7.17}$$

The same argument applies if we then apply the gauge transformation

$$\mathcal{S}_{\eta,x_0\eta/x_1,x_2;q;p}(0, ds+(d'-r_1)f-r_2e_2) \cong \mathcal{S}'_{\eta,x_0\eta/px_1x_2,\eta/px_2;q;p}(0, ds+(d+d'-r_2)f-(d-r_2)e_2), \tag{7.18}$$

and a final gauge transformation relates this to

$$\mathcal{S}'_{\eta,x_0\eta/x_1x_2,\eta/x_2,\eta/x_1,x_3,\dots,x_m;q;C}. \tag{7.19}$$

In other words, the statement holds for the generic fiber, and (tracking the various gauge transformations above) the isomorphism is given by the gauge transformation

$$\begin{aligned}\mathcal{D} \mapsto & \frac{\Gamma_{p,q}(q^{d'-d-1}z/x_0, q^{d'-2d}\eta/x_0 z, q^{r_1+r_2+d+1-d'}x_0 z/x_1x_2, q^{r_1+r_2+2-d'}x_0\eta/zx_1x_2)}{\Gamma_{p,q}(q^{r_1}z/x_1, q^{r_1+1-d}\eta/zx_1, q^{r_2}z/x_2, q^{r_2+1-d}\eta/zx_2)} \\ & \times \mathcal{D}\frac{\Gamma_{p,q}(z/x_1, q\eta/zx_1, z/x_2, q\eta/zx_2)}{\Gamma_{p,q}(z/qx_0, \eta/x_0 z, qx_0 z/x_1x_2, q^2x_0\eta/zx_1x_2)}.\end{aligned} \tag{7.20}$$

It will thus suffice to give an extension of this gauge transformation to the full algebraic moduli stack. In other words, we need to construct a family of cohomologous cocycles which for analytic curves induces the above gauge transformation (possibly up to a scalar gauge transformation).

Now, at the origin, the above product of Γ functions satisfies the q-difference equation

$$\frac{F(qz)}{F(z)} = \frac{\theta_p(\eta/x_1z, \eta/x_2z, z/qx_0, qx_0z/x_1x_2)}{\theta_p(z/x_1, z/x_2, \eta/qzx_0, qx_0\eta/zx_1x_2)}. \tag{7.21}$$

But this cocycle extends in a canonical way to the moduli stack. Indeed, we can characterize the right-hand side as the unique elliptic function of x_0 with divisor $[z/q][x_1x_2/qz]/[\eta/qz][zx_1x_2/q\eta]$ that takes the value 1 when $x_0 = x_1/q$. This gives us the required family of cocycles, and it remains to show that the various cocycles are cohomologous on the full stack. But this follows immediately from the fact that the functions representing the cocycles are elliptic in all parameters, a trivial consequence of the extension to the full stack. □

The morphism of Proposition 6.15 extends similarly.

Proposition 7.5. *Suppose $q \in \mathrm{Pic}^0(C)$ is an r-torsion point, with corresponding isogeny $\phi : C \to C'$. Then there is a functor*

$$\Phi : \mathcal{S}'_{\phi_*\circ\rho;1;C'} \to \mathcal{S}'_{\rho;q;C} \tag{7.22}$$

acting on objects as $D \mapsto rD$ and on (families of) morphisms as

$$\sum_{0\le k\le d} c_k T^k \mapsto \sum_{0\le k\le d} \phi^* c_k T^{rk}. \tag{7.23}$$

The remaining symmetry involves the Fourier transform, and is thus somewhat difficult to define on the full space of difference operators. Here the fact that we have some freedom to choose the ambient sheaf is quite helpful.

Proposition 7.6. *There is an isomorphism of sheaves of $\mathbb{Z}^{m+2}$-algebras*

$$\mathcal{S}'_{\eta,x_0,x_1,x_2,\dots,x_m;q;C} \cong \mathcal{S}'_{x_0\eta/x_1,x_1,x_0,x_2,\dots,x_m;q;C} \tag{7.24}$$

which on objects acts as

$$ds + d'f - r_1e_1 - r_2e_2 - \cdots - r_me_m \mapsto (d'-r_1)s + d'f - (d'-d)e_1 - r_2e_2 - \cdots - r_me_m. \tag{7.25}$$

Proof. Since $\mathcal{S}'_{\eta,x_0,x_1;q;C}$ is a (co)flat family and its generic fiber contains the generic fiber of

$$\mathcal{S}'_{\eta,x_0,x_1,x_2,\dots,x_m;q;C}, \tag{7.26}$$

we can use it as the ambient sheaf in constructing the canonical extension. It thus remains to show that the isomorphism

$$\mathcal{S}'_{\eta,x_0,x_1;q;C} \cong \mathcal{S}'_{x_0\eta/x_1,x_1,x_0;q;C} \tag{7.27}$$

induces the isomorphism we desire in the generic case.

Again, we may reduce understanding the generic case to understanding the analytic case, where an elementary transformation maps the problem to

$$\begin{aligned}&\mathcal{S}_{\eta,x_0\eta/x_1,\eta/px_1;q;p}(0, ds + (d + d' - r_1)f - (d - r_1)e_1)\\ &\cap \bigcap_{2\leq i\leq m} \mathcal{S}_{\eta,x_0\eta/x_1,x_i;q;p}(0, ds + (d + d' - r_1)f - r_i e_1),\end{aligned} \tag{7.28}$$

where now the isomorphism is just the Fourier transform. Since the Fourier transform preserves each subspace being intersected, the claim follows immediately. □

As in the commutative case, the above symmetries generate a Coxeter group of type $W(E_{m+1})$, with simple roots (relative to the intersection form)

$$s - e_1, f - e_1 - e_2, e_1 - e_2, \ldots, e_{m-1} - e_m \tag{7.29}$$

for $\mathcal{S}'$, and

$$s - f, f - e_1 - e_2, e_1 - e_2, \ldots, e_{m-1} - e_m \tag{7.30}$$

for $\mathcal{S}$, acting in the appropriate way ($\rho \mapsto \rho \circ s_i$) on parameters. (The Coxeter relations for the action on the category are mostly straightforward; the only tricky ones can be dealt with for $m = 2$, and follow from the generic presentation given in Appendix B below.) One key observation here is that since each simple reflection corresponds to an isomorphism of appropriately chosen ambient sheaves, the action of $W(E_{m+1})$ preserves coflatness. As a result, when proving coflatness, we may restrict our attention to Hom spaces of degree D where $D \cdot s \geq 0$ for every simple root s (essentially a fundamental chamber for $W(E_{m+1})$). Indeed, per the proof of [47, Thm. 3.4], every orbit either meets this fundamental chamber or contains an element with $D \cdot f < 0$; in the latter case, the sheaf is identically 0, so automatically coflat. We may also assume $D \cdot e_m \geq 0$, since otherwise we have an isomorphism (of sheaves)

$$\mathcal{S}'_{\rho;q;C}(0, D) \cong \mathcal{S}'_{\rho;q;C}(0, D - e_m) \tag{7.31}$$

as this is true generically. Moreover, we may either assume $m \leq 1$ (where coflatness is immediate) or that $D \cdot e_m > 0$, since when $D \cdot e_m = 0$ we have an isomorphism

$$\mathcal{S}'_{\eta,x_0,\ldots,x_m;q;C}(0, D) \cong \mathcal{S}'_{\eta,x_0,\ldots,x_{m-1};q;C}(0, D). \tag{7.32}$$

To proceed further, we need information about the leading coefficients. The analogue of the canonical divisor will play a prominent role, naturally; more generally, we will also need to consider divisors corresponding to pullbacks of the anticanonical curve. With this in mind, define divisors (relative to F_1)

$$C_l = 2s + 3f - e_1 - \cdots - e_l, \tag{7.33}$$

for $0 \leq l \leq m$. Note that for $m \geq 2$, a divisor with $D \cdot e_m \geq 0$ is in the fundamental chamber iff it is a nonnegative linear combination of the divisors

$$f, s + f, s + 2f - e_1, C_2, \ldots, C_m. \tag{7.34}$$

(We have imposed $m + 2$ inequalities on the element $D \in \mathbb{Z}^{m+2}$, thus cutting out a simplicial cone; since the cone is unimodular, the integral points are generated by the $m + 2$ divisor classes

corresponding to the dual basis.) We also define, for any divisor class $D \in \langle s, f, e_1, \ldots, e_m \rangle$ a divisor $\mathfrak{D}_{\rho;q;C}(D)$ (depending on the parameters) by

$$\begin{aligned}&\mathfrak{D}_{\rho;q;C}(ds + d'f - r_1 e_1 - \cdots - r_m e_m) \\ &\quad := \prod_{0 \le k < d} R_{q^{-k}\eta} \prod_{0 \le k < d'-d} [q^{1-k} x_0] \prod_{0 \le k < d'-2d} [\eta / q^{-k} x_0] \prod_{\substack{1 \le i \le m \\ 0 \le k < r_i}} [q^{-k} x_i]^{-1},\end{aligned} \tag{7.35}$$

extended in the natural way to products of negative length. Note that

$$\mathcal{L}(\mathfrak{D}_{\rho;q;C}(D)) \sim q^{-D \cdot (D + C_m)/2} \rho(D) \tag{7.36}$$

in $\mathrm{Pic}(C)$, so in particular has degree $C_m \cdot D$.

Lemma 7.7. *For any divisor class $D \in \langle s, f, e_1, \ldots, e_m \rangle$,*

$$T \in \mathcal{S}'_{\rho;q;C}(D - C_m, D), \tag{7.37}$$

and this induces an exact sequence

$$0 \longrightarrow \mathcal{S}'_{\rho;q;C}(0, D - C_m) \xrightarrow{T} \mathcal{S}'_{\rho;q;C}(0, D) \xrightarrow{[T^0]} \Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))), \tag{7.38}$$

where $\Gamma(C; \cdot)$ denotes the direct image to the moduli stack.

Proof. By the twisting symmetry, the first claim reduces to the claim that

$$T \in \mathcal{S}'_{\rho;q;C}(0, C_m), \tag{7.39}$$

and thus in turn to the claim that this holds generically. But this follows immediately from the fact that it holds for $m = 1$. Injectivity is then immediate from the fact that T is not a zero-divisor.

Since $\Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D)))$ is a coflat subsheaf of the flat sheaf $k(C)$, we may reduce the claim that $[T^0]\mathcal{D} \in \Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D)))$ for any family $\mathcal{D} \in \mathcal{S}'_{\rho;q;C}(0, D)$ to the generic case, where it follows immediately from the case $m = 1$.

It remains to show exactness in the middle. For this, we need to show that if $\mathcal{D}$ is a section of $\mathcal{S}'_{\rho;q;C}(0, D)$ on some open subset of parameter space, and $[T^0]\mathcal{D}$ is identically 0, then $T^{-1}\mathcal{D}$ is a section of $\mathcal{S}'_{\rho;q;C}(0, D - C_m)$. But this again reduces to the generic case. □

Remark. It will follow from the proof of coflatness that this sequence is actually exact on fibers (though the reader should note that if $\mathfrak{D}_{\rho;q;C}(D)$ is a nonconstant family of degree 0 divisors, then the fibers of $\Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D)))$ are 0, even for parameters where $\mathfrak{D}_{\rho;q;C}(D)$ is principal.)

To finish proving coflatness, it will suffice to show that this becomes a short exact sequence whenever D is in the fundamental chamber and $D \cdot e_m > 0$; we can then deduce coflatness from coflatness in degree $D - C_m$. If $D \cdot C_m < 0$ or $D \cdot C_m = 0$ with $D \neq 0$, then the leading coefficient line bundle generically has no global sections, and thus short exactness is immediate.

Lemma 7.8. *The above sequence is short exact for $D = C_m + af$ for $m \ge 2$, $a \ge 0$.*

Proof. The (generic) condition for an element of

$$\mathcal{S}'_{\eta, x_0;q;C}(0, 2s + (a+3)f) \tag{7.40}$$

to lie in

$$\mathcal{S}'_{\eta, x_0, \ldots, x_m;q;C}(0, C_m + af) \tag{7.41}$$

is precisely that its leading coefficient lies in

$$\Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(C_m + af))). \tag{7.42}$$

Since the leading coefficient sequence is short exact for $m = 0$, the claim follows. □

Note that the leading coefficient bundle has degree $D \cdot C_m$.

Lemma 7.9. *If $D_0 \in C_m + \mathbb{N}\langle f, s+f, s+2f-e_1, C_2, \dots, C_m\rangle$, then the sequence* (7.38) *is short exact for $D = D_0 + af$, $a \gg 0$.*

Proof. Suppose by induction that the claim holds for $D_0 - C_m$ (this will eventually either lead to a case with smaller m, and thus eventually to the case $m \leq 1$, or to the case $D_0 = C_m + af$ that we have already considered). Then there exists $a_1 \geq 0$ such that the sequence is short exact for $D_0 + a_1 f - C_m$. We may further assume that $(D_0 + a_1 f - C_m) \cdot C_m \geq 3$, since this condition is monotonic in a_1. Similarly, take any a_2 such that $a_2 \geq 0$ and $(C_m + a_2 f) \cdot C_m \geq 3$. Then both

$$\mathcal{S}'_{\rho;q;C}(0, D_0 + a_1 f - C_m) \qquad \text{and} \qquad \mathcal{S}'_{\rho;q;C}(D_0 + a_1 f - C_m, D_0 + (a_1 + a_2) f) \tag{7.43}$$

have precisely the expected leading coefficients, and in each case they are global sections of line bundles of degree ≥ 3. It then follows from Lemma 2.3 that the composition has the expected leading coefficients. □

It follows from this that $\mathcal{S}'(0, D)$ consists of global sections of a suitable vector bundle on $\mathbb{P}^1$. As a result, in order to prove short exactness for leading coefficients, it suffices to prove that the vector bundle is acyclic whenever

$$D \in \mathbb{N}\langle f, s+f, s+2f-e_1, C_2, \dots, C_m\rangle \tag{7.44}$$

with $D \cdot C_m > 0$, $m \geq 2$. This is certainly enough to make the leading coefficient line bundle acyclic, so we need only check that $D - C_m$ satisfies the same conditions (or has $m \leq 1$). If $m \geq 8$, this is immediate:

$$(D - C_m) \cdot C_m = D \cdot C_m + m - 8 \geq D \cdot C_m > 0. \tag{7.45}$$

while if $m < 8$, then every nonzero element of the cone satisfies $D \cdot C_m > 0$. Since the bundle is acyclic for $m \leq 1$ (so in particular for $D = 0$), we find that the bundle is always acyclic in the given cone. Combining this with the above arguments finishes our proof of the following fact.

Theorem 7.10. *Let $D, D' \in \langle s, f, e_1, \dots, e_m\rangle$. Then the* Hom *sheaf*

$$\mathcal{S}'_{\rho;q;C}(D, D') \tag{7.46}$$

is a coflat subsheaf of $\mathcal{E}ll\mathcal{D}iff(D, D')$. In particular, the resulting family of $\mathbb{Z}^{m+2}$-algebras is flat, and every fiber has a faithful representation in which the morphisms are difference operators.

Having shown this, we may now feel free to let $\mathcal{S}'_{\rho;q;C}$ denote an individual fiber of this family. Note that since this is a fiber of a family over the appropriate moduli stack, this construction is automatically functorial; in addition, all of the symmetries shown above carry over to the individual fibers. In particular, if $w \in W(E_{m+1})$, then there is an isomorphism $\mathcal{S}'_{\rho\circ w;q;C} \cong \mathcal{S}'_{\rho;q;C}$ acting as w on objects, and similarly for twisting and the adjoint.

Corollary 7.11. *For any parameters, the $\mathbb{Z}^{m+2}$-algebra $\mathcal{S}'_{\rho;q;C}$ is a domain.*

Proof. This follows immediately from the fact that the algebra of elliptic difference operators is a domain. □

We close with an application of the above ideas that does not depend on the later sheaf structure. Call a divisor class D a "rational pencil" if it is the class of a rational ruling on the generic rational surface; note that this holds iff D is in the $W(E_{m+1})$-orbit of f; similarly call a divisor class e a "-1-class" if it is in the $W(E_{m+1})$-orbit of s, or equivalently is the class of a -1-curve on the generic rational surface.

Now, suppose we are working over $\mathbb{C}$ (with $C = \mathbb{C}^*/\langle p \rangle$), and let D be a rational pencil other than f. The leading coefficient map is injective for rational pencils (as this property is preserved under $W(E_{m+1})$), so in particular every element of $\mathcal{S}'_{\rho;q;p}(0, D)$ is a difference operator of order precisely $d = D \cdot f$, so has a d-dimensional kernel (as a vector space over the field of q-elliptic functions). Call such a space a "generalized eigenspace" for D, and the nonzero elements of that space "generalized eigenfunctions". One prominent example of such generalized eigenfunctions are the elliptic biorthogonal functions of [61], which are generalized eigenfunctions of a pencil of class $2s + 2f - e_1 - e_2 - e_3 - e_4$.

One natural question in this setting is how the generalized eigenfunctions transform under various operators: if we multiply by a section of $\mathcal{S}'_{\rho;q;p}(0, f)$ (or apply some more complicated difference operator), can the result be expanded in a finite sum of generalized eigenfunctions for $\mathcal{S}'_{\rho;q;p}(f, D + f)$, and if so, how many?

To control the number of terms in such a decomposition, we may use the following result.

Lemma 7.12. *Let D be a rational pencil. For any $r > 0$, the solution space of a generic element of $\mathcal{S}'_{\rho;q;p}(0, rD)$ is a direct sum of r generalized eigenspaces for D.*

Proof. By dimension considerations, it will suffice to show that the solution space generically contains r distinct generalized eigenspaces. In particular, it will suffice to show that the generic element of $\mathcal{S}'_{\rho;q;p}(0, rD)$ admits r distinct right factors in $\mathcal{S}'_{\rho;q;p}(0, D)$. But this property is $W(E_{m+1})$-invariant (since it no longer refers to solutions), and thus we may reduce to the case $D = f$, where it is just factorization of polynomials on $\mathbb{P}^1$. □

Remark. Note that the condition on the section of $\mathcal{S}'_{\rho;q;p}(0, rD)$ is simply that the divisor of its leading coefficient (relative to the leading coefficient bundle) be multiplicity free. (More precisely, the divisor is necessarily a sum of effective divisors linearly equivalent to $\mathfrak{D}_{\rho;q;p}(D)$, and thus is invariant under a suitable hyperelliptic involution; we simply need the orbits to be distinct.)

Now, let $\mathcal{D} \in \mathcal{S}'_{\rho;q;p}(0, D)$, and let $\mathcal{D}' \in \mathcal{S}'_{\rho;q;p}(0, D')$ be another operator, with the problem being to understand how $\mathcal{D}'g$ expands in generalized eigenfunctions of $\mathcal{S}'_{\rho;q;p}(D', D' + D)$ given that $\mathcal{D}g = 0$. To apply the previous lemma, we need to find a section of $\mathcal{S}'_{\rho;q;p}(D', D' + rD)$ which when multiplied by $\mathcal{D}'$ admits $\mathcal{D}$ as a right factor. Note that if such a section is to exist for generic $\mathcal{D}'$, there should be a section that works for *all* $\mathcal{D}'$. (Indeed, the span over all $\mathcal{D}'$ should still lie in a finite union of eigenspaces, so we just need to take a suitable gcd.)

To do this, it will be convenient to blow up our surface further. Let x_{m+1} be a point where the leading coefficient of $\mathcal{D}$ vanishes as a section of the leading coefficient bundle (one of two such points). Then we have $\mathcal{D} \in \mathcal{S}'_{\rho,x_{m+1};q;p}(0, D - e_{m+1})$, where now $D - e_{m+1}$ is a -1-class. For any r, there is a unique (modulo scalars) section of

$$\mathcal{S}'_{\rho,x_{m+1};q;p}(D', D' + r(D - e_{m+1})) \subset \mathcal{S}'_{\rho,x_{m+1};q;p}(D', D' + rD), \tag{7.47}$$

giving us a natural candidate for the operator to apply. For this to work, we need any product of that section with an operator of degree D' to have $\mathcal{D}$ as a right factor. To accomplish this, we may use the following result.

Lemma 7.13. *Let e be a -1-class. If $D \cdot e < 0$, then the composition morphism*

$$\mathcal{S}'_{\rho;q;C}(e, D) \circ \mathcal{S}'_{\rho;q;C}(0, e) \to \mathcal{S}'_{\rho;q;C}(0, D) \tag{7.48}$$

is surjective.

Proof. This is a $W(E_{m+1})$-invariant statement, so we may reduce to the case $e = e_m$, when it corresponds to the fact that

$$\mathcal{S}'_{\rho;q;C}(0, D - e_m) = \mathcal{S}'_{\rho;q;C}(0, D) \tag{7.49}$$

whenever the coefficient of e_m in D is positive. □

In particular, we need simply choose r above so that $(D' + r(D - e_{m+1})) \cdot (D - e_{m+1}) < 0$, or equivalently, $r \geq D \cdot D' + 1$. This gives us the following result.

Proposition 7.14. *Let D be a rational pencil other than f, and let D' be any other divisor class. Let g be a generic generalized eigenfunction for D. Then for any $\mathcal{D} \in \mathcal{S}'_{\rho;q;p}(0, D')$, $\mathcal{D}g$ can be expressed as a sum of at most $D \cdot D' + 1$ generalized eigenfunctions for $\mathcal{S}'_{\rho;q;p}(D', D' + D)$.*

We can be more precise, of course: the generalized eigenvalues can be labelled by pairs of points on $\mathbb{C}^*/\langle p \rangle$ (i.e., the two points where the leading coefficient of the operator annihilating g vanishes), and the new generalized eigenvalues correspond to points of the form $q^{-k}x$, $0 \leq k \leq D \cdot D'$, where x is one of the original two points. The genericity condition is that these form $D \cdot D' + 1$ distinct pairs under the appropriate involution.

As an example, consider the case $D = 2s + 2f - e_1 - e_2 - e_3 - e_4$, $D' = f$. In this case, the claim is that multiplying a generalized eigenfunction by any section of $\mathcal{S}'_{\rho;q;C}(0, f)$ gives a linear combination of three generalized eigenfunctions (which must be elliptic functions if the original was an elliptic function); the resulting (known) expansion is essentially the analogue for the elliptic biorthogonal functions of the three-term recurrence for orthogonal polynomials. Similarly, the fact that the unique operator of degree $D' = s + f - e_1 - e_2$ takes biorthogonal functions to biorthogonal functions with shifted parameters corresponds to the fact that $D \cdot D' = 0$ in that case, and similarly for the lowering operator $D' = s$ and the raising operator $D' = s + 2f - e_1 - e_2 - e_3 - e_4$. The latter case gives rise to Rodrigues-type formulas: for any $r \geq 0$, the image of 1 under the unique section of $\mathcal{S}'_{\rho;q;p}(-r(s + 2f - e_1 - e_2 - e_3 - e_4), 0)$ is a generalized eigenfunction, and the operator of that degree can be obtained from the operator of degree rs (roughly speaking, the rth power of the lowering operator) via gauging by a product of Γ functions. Another example is the case $D' = e_4$, with $D \cdot D' = 1$; here the operator itself is trivial, but the equations change, giving an instance of a sparse change of basis matrix between two families of biorthogonal functions.

8 Surjectivity of composition

Of course, the above flat family of $\mathbb{Z}^{m+2}$-algebras is not quite a family of noncommutative surfaces, for the simple reason that they do not come with a natural notion of sheaf. As we will see below, there are several natural notions of sheaf (or, rather, of "torsion"), and it is nontrivial to show that they all agree. In order to prove this, we will need a number of results about surjectivity of composition morphisms. The results we obtain along these lines will all hold subject to conditions

on ρ that are invariant under the twist isomorphism, and thus it will always suffice to show them in the case that the domain of the composed morphism is the 0 object. With this in mind, we introduce the following short-hand notation:

$$[D_1][D_2]\cdots[D_n] \tag{8.1}$$

denotes the image of the composition morphism

$$\begin{aligned}\mathcal{S}'_{\rho;q;C}(D_2+\cdots+D_n, D_1+\cdots+D_n)\otimes\cdots\otimes\mathcal{S}'_{\rho;q;C}(D_n, D_{n-1}+D_n)\otimes\mathcal{S}'_{\rho;q;C}(0,D_n) \\ \subset \mathcal{S}'_{\rho;q;C}(0, D_1+\cdots+D_n).\end{aligned} \tag{8.2}$$

Thus for instance

$$[D_1][D_2]\cdots[D_n] = [D_1+\cdots+D_n] \tag{8.3}$$

iff the above composition morphism is surjective, and to show that $[D_1][D_2][D_3] = [D_1+D_2+D_3]$ for a twist-invariant subset of parameters, it would suffice to show that $[D_1][D_2] = [D_1+D_2]$ and $[D_1+D_2][D_3] = [D_1+D_2+D_3]$.

Most of the results we will consider will involve divisor classes in the following cone. Say that D is "universally nef" if $D \in \mathbb{N}\langle f, s+f, s+2f-e_1, C_2, \ldots, C_m\rangle$ and $D\cdot C_m \geq 0$; we will call the corresponding convex cone the "universal nef cone". (These are the divisor classes on m-fold blowups of commutative F_1 which are nef on all such blowups admitting smooth anticanonical curves.) Note that if $m \leq 8$, then the additional constraint is redundant, since all of the generators of the simplicial cone have $D\cdot C_m \geq 0$. Also, given any divisor class

$$D = af + b(s+f) + c(s+2f-e_1) + d_2C_2 + \cdots d_mC_m, \tag{8.4}$$

define

$$X_l(D) := af + b(s+f) + c(s+2f-e_1) + \sum_{2\leq i\leq l} d_iC_i \tag{8.5}$$

for $1 \leq l \leq m$, as well as

$$X'_0(D) := af + b(s+f) \qquad X_0(D) := af + c(s+2f-e_1); \tag{8.6}$$

note that $X_l(D)$ is supported on the l-fold blowup of F_1, while $X'_0(D)$ is supported on F_1 and $X_0(D)$ is supported on F_0/F_2.

Lemma 8.1. *Let D be a universally nef divisor class with $D\cdot C_m > 0$, D not a multiple of f. Then the leading coefficient morphism*

$$[D] \to \Gamma(\mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))) \tag{8.7}$$

is surjective.

Proof. This follows from known results for $m \leq 1$, so we may as well assume $m \geq 2$. We may also assume that $D\cdot e_m \neq 0$, since otherwise $D = X_{m-1}(D)$ and we may as well work on in the corresponding subcategory. It follows that $D - C_m$ is universally nef (it is certainly in the relevant simplicial cone, so this is immediate if $m \leq 8$, and for $m > 8$ we have $(D-C_m)\cdot C_m > D\cdot C_m$). We may thus use flatness to compute

$$\dim[D] - \dim[D-C_m] = D\cdot C_m = \dim\Gamma(\mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))). \tag{8.8}$$

□

Remark. Since the surjectivity of the leading coefficient morphism is invariant under $W(E_{m+1})$, this actually applies to any divisor class which is $W(E_{m+1})$-equivalent to a class of the above form; in other words $W(E_{m+1})D$ meets the universal nef cone, $D \cdot C_m > 0$, and D is not a multiple of a rational pencil.

It turns out that this is enough to mostly let us reduce to the case $m \leq 7$.

Lemma 8.2. *Suppose D_1, D_2 are universally nef divisor classes with $D_i \cdot C_m \geq 3$ such that $[X_7(D_1)][X_7(D_2)] = [X_7(D_1 + D_2)]$ for all ρ, q, C. Then $[D_1][D_2] = [D_1 + D_2]$ for all ρ, q, C.*

Proof. If $m \leq 7$, the claim is immediate; if $D_1 \cdot e_m = D_2 \cdot e_m = 0$, then we may work entirely in the blown down subcategory, and thus the claim will hold by induction. We may thus assume (using the adjoint symmetry as necessary) that $m > 7$ and $D_1 \cdot e_m > 0$. It follows that $D_1 - C_m$ is again universally nef, and since $m > 7$,

$$(D_1 - C_m) \cdot C_m \geq D_1 \cdot C_m \geq 3. \tag{8.9}$$

It follows, therefore (by induction on $D_1 \cdot e_m$) that $[D_1 - C_m][D_2] = [D_1 + D_2 - C_m]$. In particular, we find that

$$T[D_1 + D_2 - C_m] = T[D_1 - C_m][D_2] \subset [D_1][D_2], \tag{8.10}$$

and thus to prove surjectivity, it remains only to prove it for the leading coefficients. But this follows immediately from Lemma 2.3. □

Remark. Note that since the constraints $D_i \cdot C_m \geq 3$ only arise in applying Lemma 2.3, the argument works equally well if $D_1 \cdot C_m = 2$, $D_2 \cdot C_m \geq 3$, or vice versa.

Here and in the remainder of the section, the phrase "for all ρ, q, C" should be understood, in the absence of a specific condition otherwise.

Lemma 8.3. *Suppose D_1, D_2 are universally nef divisor classes with $X_5(D_1) = D_1$, $X_5(D_2) = D_2$. Then $[D_1][D_2] = [D_1 + D_2]$.*

Proof. We first note that this holds if $D_1 \in \{C_1, \ldots, C_5\}$. Indeed, in that case, we find that $[T^0][D_1]$ are the global sections of a line bundle of degree ≥ 3, and thus for any operator in $[D_1 + D_2]$, there is a corresponding sum of compositions with the same leading coefficient. Since $T[D_2] \subset [C_i][D_2]$, every operator in $[D_1 + D_2]$ with vanishing leading coefficient is in $[C_i][D_2]$, and the claim follows. Taking adjoints, we also find that the claim holds if $D_2 \in \{C_1, \ldots, C_5\}$.

Now, suppose more generally that $D_1 - C_i$ is universally nef for some i. Then we have

$$[D_1][D_2] = [D_1 - C_i][C_i][D_2] = [D_1 - C_i][D_2][C_i]. \tag{8.11}$$

In particular, if $[D_1 - C_i][D_2] = [D_1 + D_2 - C_i]$, then $[D_1][D_2] = [D_1 + D_2]$; as a result (and using the adjoint symmetry), we may reduce to the case that none of $D_i - C_j$ are universally nef.

It follows in particular that each of D_1, D_2 is in one of the simplicial cones $\mathbb{N}\langle f, s + f\rangle$, $\mathbb{N}\langle f, s + 2f - e_1\rangle$. If they are in the same cone, then we reduce to known surjectivity results for F_1 or F_0 as appropriate. Thus suppose that D_1 is in the first cone (but not the second) and D_2 is in the second cone (but not the first); the other case will follow by taking adjoints. We then have

$$\begin{aligned}[D_1][D_2] &= [D_1 - s - f][s + f][s + 2f - e_1][D_2 - s - 2f + e_1] \\ &= [D_1 - s - f][C_1][D_2 - s - 2f + e_1] \\ &= [D_1 - s - f][D_2 - s - 2f + e_1][C_1],\end{aligned} \tag{8.12}$$

and the claim again follows by induction. Here the fact that $[s + f][s + 2f - e_1] = [2s + 3f - e_1]$ is just Lemma 3.3 combined with surjectivity for leading coefficients. □

In this strong form, this fails for $m > 5$; in particular, C_6 is universally nef, but $[C_6][C_6] \subsetneq [2C_6]$ when $q^2 = 1$. Indeed, the leading coefficients of the two copies of $[C_6]$ then lie in a degree 2 line bundle, and thus we only obtain three of the four requisite leading coefficients of $[2C_6]$. A related issue is that there are parameters for which $[f][C_6] \subsetneq [f + C_6]$, since it is possible for the two leading coefficient bundles to be isomorphic. These are essentially the only issues for $m = 6$, and indeed we have the following.

Lemma 8.4. *Let $D = X_7(D)$ be a divisor class other than C_6 such that $X_6(D)$ is not a multiple of f. Then $[C_6][D] = [D][C_6] = [D + C_6]$.*

Proof. If $X_6(D) \cdot C_6 \geq 3$, then we can argue as before: we have surjectivity on leading coefficients, and either $D \neq X_6(D)$, so we reduce to the $D - C_7$ case, or $D = X_6(D)$, in which case usual argument applies, using $T \in [C_6]$. (Note that this includes the case $D = 2C_6$.)

Since $X_6(D)$ is not a multiple of f, the only surviving possibility for $X_6(D)$ is C_6 itself, and the same reduction applies until we reach the case $D = C_6 + C_7$. Here, we again have surjectivity on the leading coefficients, but the usual argument would require $[C_6][C_6] = [2C_6]$. Luckily, it suffices to show that

$$[C_6 - C_7][C_6 + C_7] + [C_6][C_6] = [2C_6]. \tag{8.13}$$

This actually *is* surjective on leading coefficients; the first term gives a codimension 1 subspace of $[T^0][2C_6]$, consisting of sections vanishing at the appropriate twist of x_7, and it is straightforward to check that $[C_6][C_6]$ contains a leading coefficient not in that subspace. Since $[C_6 - C_7][C_6] = [e_7][C_6] = [C_6 + e_7]$, the claim follows in general. □

Since $C_7^2 = 1$, multiplication by $[C_7]$ cannot be surjective on leading coefficients, so we must consider multiples thereof.

Lemma 8.5. *Let the universally nef divisor class $D = X_7(D)$ be such that $D \cdot C_7 \geq 3$ and $D \cdot e_7 > 0$. Then for all $r > 1$, $[rC_7][D] = [D][rC_7] = [D + rC_7]$.*

Proof. The usual argument reduces this to the case $r = 2$, and since we have surjectivity on leading coefficients, to showing that

$$[C_7][D] + [2C_7][D - C_7] = [D + C_7]. \tag{8.14}$$

As in the previous Lemma, it suffices to show surjectivity of leading coefficients, the first term gives a codimension 1 subspace vanishing at a particular point, and the second term includes leading coefficients not vanishing there. □

Remark. The condition $D \cdot e_7 > 0$ is necessary for the argument, as otherwise the second term *also* has a nontrivial gcd, which will agree with the gcd of the first term on a nonempty closed substack of parameter space.

Lemma 8.6. *Let D be a universally nef divisor class such that $D \cdot C_m \geq 3$ and $X_5(D) \neq 0$. Then $[f][D] = [D][f] = [D + f]$.*

Proof. We may as usual assume that $D \cdot e_m > 0$, and that $m \in \{6, 7\}$ (since the case $m \leq 5$ has already been shown). Surjectivity on leading coefficients is immediate from Lemma 2.3, and $[f][D - C_m] = [f + D - C_m]$ by induction on $D \cdot (e_6 + e_7)$. □

Proposition 8.7. *Let D_1, D_2 be universally nef divisor classes such that $D_i \cdot C_m \geq 3$ and both $X_5(D_1)$ and $X_5(D_2)$ are nonzero. Then $[D_1][D_2] = [D_1 + D_2]$.*

Proof. We may assume $m \leq 7$, and by the usual argument that $D_1 \cdot e_m > 0$; moreover, if $m \leq 5$, we have already shown this to be true. Since we automatically have surjectivity on leading coefficients, we may proceed by induction, with the one exception being when $(D_1 - C_m) \cdot C_m < 3$. By inspection of the generators of the (simplicial) universal nef cone, this implies that $D_1 - C_m \in \{0, f, C_6, C_7, 2C_7\}$; together with the assumption $X_5(D_1) \neq 0$, this yields the possibilities $D_1 = C_6 + f$ for $m = 6$ and $D_1 = C_7 + f$ for $m = 7$. But in that case, we can use the previous lemma as an additional base case for the induction. □

Our most general surjectivity result is the following.

Proposition 8.8. *Let D_1, D_2 be universally nef divisor classes with $D_1 \cdot C_m \geq 2$, $D_2 \cdot C_m \geq 3$, and suppose that D_1 is contained in the closure of the open face of the universal nef cone containing D_2. Then $[D_1][D_2] = [D_2][D_1] = [D_1 + D_2]$.*

Proof. The usual reduction applies to let us assume $m \in \{6, 7\}$. If $D_2 \cdot e_m = 0$, then $D_1 \cdot e_m = 0$ (every bounding inequality of the universal nef cone which is tight for D_2 must by assumption be tight for D_1 as well), so that we may as well assume $D_2 \cdot e_m > 0$. Similarly, if $X_5(D_1) \neq 0$, then $X_5(D_2) \neq 0$, and again we have already shown surjectivity of multiplication.

If $m = 6$, then $D_1 \propto C_6$, while the conditions on D_2 imply $D_2 \neq C_6$ and $D_2 \not\propto f$, and thus the claim reduces to the corresponding claim for $D_1 = C_6$, shown above.

The remaining case case $m = 7$, $D_1 = aC_6 + bC_7$. If $a > 0$ and $X_6(D_2) \cdot C_6 \geq 3$, then we may reduce to the $m = 6$ case in the usual way. Since $X_5(D_2) \neq X_6(D_2)$, the only other case for $a > 0$ is that $X_6(D_2) = C_6$. Since $D_2 \neq C_6$, we can again reduce to the case $D_1 = C_6$.

Finally, if $a = 0$, so $D_1 \propto C_7$, then the claim is just Lemma 8.5. □

Applying this to the case $D_2 \propto D_1$ gives the following.

Corollary 8.9. *If D is a universally nef divisor class with $D \cdot C_m \geq 3$, then the sub-$\mathbb{Z}$-algebra with objects $\mathbb{Z}D$ is generated in degree 1.*

Also of interest are the $\mathbb{Z}$-algebras corresponding to the anticanonical class for $m \in \{6, 7\}$. Note that we can use the elements $T \in [C_6]$ to turn these $\mathbb{Z}$-algebras into graded algebras, which are then generated in the specified degrees. (For $m \leq 5$, they are of course generated in degree 1.)

Corollary 8.10. *The sub-$\mathbb{Z}$-algebra with objects $\mathbb{Z}C_6$ is locally generated by morphisms of degrees 1 and 2; the sub-$\mathbb{Z}$-algebra with objects $\mathbb{Z}C_7$ is locally generated by morphisms of degrees 1, 2, 3.*

Proof. The only claim that does not immediately follow from the surjectivity results above is that $[4C_7]$ can be expressed in terms of lower degrees, or equivalently that $[4C_7] = [C_7][3C_7] + [2C_7][2C_7] + [3C_7][C_7]$. But this follows from the usual argument: the first term is codimension 1, with leading coefficients having nontrivial gcd, letting us verify that the second term is not contained in the first. □

The above results will be enough to let us identify the various notions of torsion that hold universally, but in general we expect that there should be far more ample divisors than just those coming from the interior of the universal nef cone. Indeed, we expect that every ample divisor on the generic commutative rational surface of our form should deform to an ample divisor on the generic noncommutative rational surface in our family. Relatedly, we expect the $W(E_{m+1})$ symmetry to extend (at least generically) to the categories of sheaves, which will necessitate having notions of torsion which are at least invariant under subgroups of the Weyl group.

To control this, we will need some surjectivity results that only hold away from countable unions of hypersurfaces in parameter space. We will always state the conditions in translation-invariant ways; in general, of course, this will be far more stringent than necessary (the failure to surject is a closed condition, so only finitely many of the excluded hypersurfaces actually need to be excluded), but we will have better methods below for dealing with such cases, and translation-invariance makes the arguments simpler.

Lemma 8.11. *Let $S_1, S_2 \subset \{1, \dots, m\}$ be subsets for which $\{x_i : i \in S_1 \setminus S_2\}$ and $\{x_i : i \in S_2 \setminus S_1\}$ are disjoint in $C/q^{\mathbb{Z}}$, and define $D_i = 2s + (a_i + 3)f - \sum_{j \in S_i} e_j$, where $a_i \geq \max(|S_i|/2 - 5/2, 0)$. Then $[D_1 + a_1 f][D_2 + a_2 f] = [D_1 + D_2 + (a_1 + a_2)f]$.*

Proof. The constraints on a_1, a_2 are enough to ensure that the leading coefficient bundles have degree ≥ 3, and thus we have surjectivity on leading coefficients. The claim then reduces to showing

$$[D_1 - C_m][D_2] + [D_1][D_2 - C_m] = [D_1 + D_2 - C_m] \tag{8.15}$$

Since $D_1 - C_m = a_1 f + \sum_{j \notin S_1} e_j$, we see that $[D_1 - C_m][D_2] = [\sum_{j \notin S_1} e_j][a_1 f][D_2]$ and similarly for the second term. We thus see that the first term gives the subspace of the leading coefficient bundle obtained by imposing vanishing at suitable translates of x_j for $j \in S_2$, and similarly for the second term. It then follows immediately that we again have surjectivity on leading coefficients, and thus may reduce to

$$[D_1 - C_m][D_2 - C_m] = [D_1 + D_2 - 2C_m] \tag{8.16}$$

But this reduces immediately to $[a_1 f][a_2 f] = [(a_1 + a_2)f]$, and the claim follows. □

Lemma 8.12. *If $\eta/x_1 x_2$ is not a power of q, then for $b_1 \geq 4$, $b_2 \geq 3$,*

$$[2s + b_1 f - 2e_1 - e_2][2s + b_2 f - e_1] = [4s + (b_1 + b_2)f - 3e_1 - e_2] \tag{8.17}$$

Proof. This is essentially just an elementary transformation of the same argument (only on $\mathcal{S}$ rather than $\mathcal{S}'$). □

This of course generalizes as much as Lemma 8.11, but we will only need the following somewhat technical special case, again based on $\mathcal{S}$ rather than $\mathcal{S}'$ but otherwise the same argument.

Lemma 8.13. *Suppose that the points $\eta/x_1, x_2, \dots, x_m \in C/q^{\mathbb{Z}}$ are distinct, and let $D_i = 2s + 4f - 2e_1 - \sum_{m-i<k\leq m} e_k$. Then for $1 \leq i < m$,*

$$[C_i + if][D_{m-i} + (m-i)f] = [D_{m+1-i} + (m-i)f][C_i + if] = [C_i + D_{m+1-i} + mf]. \tag{8.18}$$

Remark. Note that

$$C_i + D_{m+1-i} = C_m + 2s + (m+4)f - 2e_1, \tag{8.19}$$

which is independent of i, and is universally nef.

The above cases are related to the two subsystems A_{m-1}; there is also a surjectivity result similarly related to D_m. Let ι_η denote the (hyperelliptic) involution on C corresponding to η.

Lemma 8.14. *For $m \geq 2$, suppose that the images of $x_1, \dots, x_m$ under the quotient of C by $\langle q^{\mathbb{Z}}, \iota_\eta \rangle$ are distinct. Then for any $S \subset \{1, \dots, m\}$,*

$$\Big[s + (|S| + 1)f - \sum_{i \in S} e_i\Big]\Big[s + (m - |S| + 1)f - \sum_{i \notin S} e_i\Big] = [C_m + (m-1)f]. \tag{8.20}$$

Proof. We have surjectivity on leading coefficients since the two leading term bundles have degree $|S|+3$ and $m-|S|+3$ respectively, both ≥ 3. It thus remains only to show that we obtain all requisite multiples of T. For $j \in S$, we may write

$$\begin{aligned}
&\big[s+(|S|+1)f-\sum_{i\in S}e_i\big]\big[s+(m-|S|+1)f-\sum_{i\notin S}e_i\big] \\
&\qquad\supset \big[(|S|-1)f-\sum_{i\in S\setminus\{j\}}e_i\big][s+2f-e_j][s+f]\big[(m-|S|)f-\sum_{i\notin S}e_i\big] \\
&\qquad= \big[(|S|-1)f-\sum_{i\in S\setminus\{j\}}e_i\big][2s+3f-e_j]\big[(m-|S|)f-\sum_{i\notin S}e_i\big] \\
&\qquad\supset \big[(|S|-1)f-\sum_{i\in S\setminus\{j\}}e_i\big]\big[\sum_{1\leq i\leq m;i\neq j}e_i\big]\big[(m-|S|)f-\sum_{i\notin S}e_i\big]T.
\end{aligned}$$

Similarly, for $j \notin S$, we find

$$\begin{aligned}
&\big[(|S|)f-\sum_{i\in S}e_i\big]\big[\sum_{1\leq i\leq m}e_i\big]\big[(m-|S|-1)f-\sum_{i\notin S\cup\{j\}}e_i\big]T \\
&\qquad\subset \big[s+(|S|+1)f-\sum_{i\in S}e_i\big]\big[s+(m-|S|+1)f-\sum_{i\notin S}e_i\big].
\end{aligned}$$

We have thus obtained a total of m 1-dimensional subspaces of $[(m-1)f]$, each consisting of sections vanishing on all but one of a fixed set of m distinct points of $\mathbb{P}^1$. In particular, those subspaces span $[(m-1)f]$, and the claim follows. □

We will only need this in the following form.

Corollary 8.15. *For $m \geq 2$, suppose that the images of $x_1, \ldots, x_m$ under the quotient of C by $\langle q^{\mathbb{Z}}, \iota_\eta\rangle$ are distinct. Then*

$$[C_i+(i-1)f][2C_m-C_i+(2m-i)f] = [2C_m+(2m-2)f]. \tag{8.21}$$

Proof. We have

$$\begin{aligned}
&[C_i+(i-1)f][2C_m-C_i+(2m-i)f] \\
&\supset [s+(i+1)f-\sum_{1\leq j\leq i}e_j][s+f][s+(m+1)f-\sum_{1\leq j\leq m}e_j][s+(m-i+1)f-\sum_{i<j\leq m}e_j] \\
&= [s+(i+1)f-\sum_{1\leq j\leq i}e_j][C_m+(m-1)f][s+(m-i+1)f-\sum_{i<j\leq m}e_j] \\
&= [s+(i+1)f-\sum_{1\leq j\leq i}e_j][s+(m-i+1)f-\sum_{i<j\leq m}e_j][C_m+(m-1)f] \\
&= [C_m+(m-1)f][C_m+(m-1)f] \\
&= [2C_m+(2m-2)f] \qquad (8.22)
\end{aligned}$$

as required. □

We also recall a surjectivity result from the F_0 case (which over F_0 simply states that $[s][f] = [f][s] = [s+f]$ under the appropriate constraint).

Lemma 8.16. *If x_0/x_1 is not a power of q, then $[s+f-e_1][f] = [f][s+f-e_1] = [s+2f-e_1]$.*

We will need one more surjectivity result for the case that q is torsion, in order to show that the resulting surfaces are finite over commutative surfaces. Let q be an r-torsion point with C' the corresponding quotient of C, and recall the functor Φ of Proposition 7.5. Here we only need to prove surjectivity for a single sufficiently ample D.

Lemma 8.17. *For all $m \geq 1$ there exists $a \geq 0$ such that for $D = af + \sum_{1 \leq l \leq m} C_l$, the multiplication map*

$$\mathcal{S}'_{\rho;q;C}(rD, 2rD) \otimes \Phi(\mathcal{S}'_{\phi_* \circ \rho;1;C'}(0, D)) \to \mathcal{S}'_{\rho;q;C}(0, 2rD) \tag{8.23}$$

is surjective.

Proof. Note that for $a \gg 0$, D is in the interior of the universal nef cone. We also observe that there is an analogue of Lemma 2.3 in this situation, which states that if $\mathcal{L}$ is a line bundle on C of degree ≥ 3, then for any line bundle $\mathcal{L}'$ on C' of degree ≥ 2, the multiplication map

$$\Gamma(C; \mathcal{L}) \otimes \Gamma(C'; \mathcal{L}') \to \Gamma(C; \mathcal{L} \otimes \phi^* \mathcal{L}') \tag{8.24}$$

is surjective. (If $\mathcal{L} \otimes (\phi^* \mathcal{L}')^{-1}$ is acyclic, then we can use the same argument as in the first part of Lemma 2.3, and simply take the two relatively prime global sections of $\phi^* \mathcal{L}'$ to be pullbacks of global sections of $\mathcal{L}'$; otherwise, we may reduce to the case $\deg \mathcal{L}' = 2$.)

We may thus use this (increasing a as necessary to make the degrees sufficiently large) to give surjectivity on leading coefficients for all cases of the form $[rD - bC_m]\Phi([D])$ or $[rD - bC_m]\Phi([D - C_m])$, $0 \leq b < r$, allowing us to induct in m. The base case is $D = af$, for which surjectivity for $a \gg 0$ follows from the fact that ϕ descends to a degree r map on $\mathbb{P}^1$. □

Remark. Note that the same argument applies to $D = 2s + af$ for $a \gg 0$ in either case with $m = 0$. For the $\mathbb{P}^2$ case, we may refer to [4, §7].

9 Torsion modules and sheaves II

As we mentioned, there are several ways we could try to define "torsion" for $\mathcal{S}'_{\rho;q;C}$-modules. Since we intend to prove these are all equivalent, we may as well use the one which is simplest to work with. If v is a homogeneous element of a $\mathcal{S}'_{\rho;q;C}$-module, we denote by $[D]v$ the space of all images of v under the action of $\mathcal{S}'_{\rho;q;C}(\deg(v), \deg(v) + D)$.

Definition 9.1. Let M be a $\mathcal{S}'_{\rho;q;C}$-module. A homogeneous element $v \in M$ is *torsion* if there is some divisor class D_0 such that for any universally nef divisor class D, $[D + D_0]v = 0$. If every homogeneous element of M is torsion, we say that M is torsion.

Call a divisor class "universally very ample" if it is in the interior of the universal nef cone and satisfies $D \cdot C_m \geq 3$. This notion gives us an alternate description of when an element is torsion.

Proposition 9.1. *A homogeneous element v is torsion iff there is a universally very ample divisor class D such that $[D]v = 0$.*

Proof. One direction is easy: any translate of the universal nef cone meets the interior of the universal nef cone; if D is in the intersection, then so is $3D$, and $3D$ will be universally very ample.

Thus suppose $[D]v = 0$ with D universally very ample. Then for any universally nef divisor class D' with $D' \cdot C_m \geq 2$, Proposition 8.8 tells us that $[D'][D] = [D' + D]$, and thus $[D' + D]v = [D'][D]v = 0$. Since the space of such $D' + D$ contains a translate of the universal nef cone, v is torsion. □

Call a class "universally ample" if it is in the interior of the universal nef cone.

Proposition 9.2. *Let D be any universally ample class. Then a homogeneous element v is torsion iff there exists $a \geq 0$ such that $[bD]v = 0$ for all $b \geq a$.*

Proof. If D is universally ample, then $3D$ is universally very ample, so one direction is trivial. For the other direction, we need simply observe that for any D_0, $aD - D_0$ is universally nef for $a \gg 0$, and thus any translate of the universal nef cone contains all sufficiently large multiples of D. □

The converse can be strengthened.

Proposition 9.3. *Let D be a universally nef class with $D \cdot C_m \geq 2$. If $[D]v = 0$ for some homogeneous element v, then v is torsion.*

Proof. Let D_2 be a universally very ample class, and note that Proposition 8.8 tells us that $[D_2][D] = [D + D_2]$. But then v is annihilated by $[D + D_2]$; since $D + D_2$ is universally very ample, the claim follows. □

Proposition 9.4. *The torsion modules form a Serre subcategory of $\mathcal{S}'_{\rho;q;C} - \mathcal{M}od$.*

Proof. That "torsion" is inherited by sub- and quotient modules is trivial (indeed, we may use the same translate of the universal nef cone). Thus it remains only to consider extensions. Suppose we have a short exact sequence $0 \to M_1 \to M \to M_2 \to 0$ with M_1, M_2 torsion, and consider a homogeneous element $v \in M$. The image of v in M_2 is torsion by assumption, and thus there exists universally very ample D_1 such that $[D_1]v \in M_1$. Choose a basis of $w_1, \ldots, w_n \in [D_1]v$, and note that since M_1 is torsion, each w_i is torsion. It follows, therefore, that each w_i is annihilated by some translate of the universal nef cone. The intersection of those translates will itself contain a translate of the universal nef cone; it follows that there is a divisor D_0 such that for all universally nef D, $[D + D_0]w_i = 0$ for $1 \leq i \leq n$. It follows that $[D + D_0][D_1]v = 0$, and thus as long as $(D + D_0) \cdot C_m \geq 2$, $[D + D_0 + D_1]v = 0$. The resulting set of divisors contains a translate of the universal nef cone, and thus v is indeed torsion. □

Definition 9.2. A "quasicoherent sheaf on $X_{\rho;q;C}$" is an object in the quotient of $\mathcal{S}'_{\rho;q;C} - \mathcal{M}od$ by the Serre subcategory of torsion modules.

Note that for any divisor class D, the object $-D$ induces a projective $\mathcal{S}'_{\rho;q;C}$-module, namely the functor $\mathcal{S}'_{\rho;q;C}(-D, _)$. The corresponding sheaf will be denoted by $\mathcal{O}_X(D)$, with (D) omitted when $D = 0$; we will call these sheaves "line bundles" by obvious analogy with the commutative case. We let $X_{\rho;q;C}$ denote the quasi-scheme corresponding to the category $\operatorname{qcoh} X_{\rho;q;C}$ of quasicoherent sheaves and with marked object $\mathcal{O}_X$.

Definition 9.3. The *saturated $\mathbb{Z}^{m+2}$-algebra* $\hat{\mathcal{S}}'_{\rho;q;C}$ is defined by

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2) := \operatorname{Hom}(\mathcal{O}_X(-D_2), \mathcal{O}_X(-D_1)) \tag{9.1}$$

We will discuss this $\mathbb{Z}^{m+2}$-algebra in more detail below, but will need a few facts now. Note that since $\mathcal{S}'_{\rho;q;C}$ is a domain, the natural functor $\mathcal{S}'_{\rho;q;C} \to \hat{\mathcal{S}}'_{\rho;q;C}$ is faithful. We also have the following key point: the morphisms in the saturated $\mathbb{Z}^{m+2}$-algebra are still difference operators. More precisely, we have the following.

Lemma 9.5. *The space $\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2)$ can be identified (compatibly with composition) with the space of elliptic difference operators $\mathcal{D}$ such that*

$$\mathcal{S}'_{\rho;q;C}(D_2, D_3)\mathcal{D} \subset \mathcal{S}'_{\rho;q;C}(D_1, D_3) \tag{9.2}$$

for all D_3 in some translate of the universal nef cone.

Proof. It follows from the definition of torsion that any morphism from $\mathcal{O}_X(-D_2)$ to $\mathcal{O}_X(-D_1)$ is represented by a compatible family of maps

$$\Phi : \mathcal{S}'_{\rho;q;C}(D_2, D_3) \to \mathcal{S}'_{\rho;q;C}(D_1, D_3) \tag{9.3}$$

for D_3 ranging over a translate of the universal nef cone. We thus need to show that this family can be expressed as right multiplication by a difference operator. Let

$$D = \sum_{1 \le i \le m} (C_i + if) \tag{9.4}$$

and note that any translate of the universal nef cone contains some point of the form $D_2 + rD$. Now, $\mathcal{S}'_{\rho;q;C}(D_2, D_2 + rD)$ contains an operator of the form $g(z)T^{rm}$, and for any difference operator $\mathcal{D} \in \mathcal{S}'_{\rho;q;C}(D_2, D_3)$, there is an integer $a \ge 0$ and a morphism

$$G(z) \in \mathcal{S}'_{\rho;q;C}(D_3 + rC_m, D_3 + rD + af) \tag{9.5}$$

such that

$$G(z)T^{rm}\mathcal{D} \in \mathcal{S}'_{\rho;q;C}(D_2 + rD, D_3 + rD + af)g(z)T^{rm}. \tag{9.6}$$

Then

$$\begin{aligned} \Phi(\mathcal{D}) &= T^{-rm}G(z)^{-1}\Phi(G(z)T^{rm}\mathcal{D}) \\ &= T^{-rm}G(z)^{-1}\Phi(G(z)T^{rm}\mathcal{D}T^{-rm}g(z)^{-1}g(z)T^{rm}) \\ &= \mathcal{D}T^{-rm}g(z)^{-1}\Phi(g(z)T^{rm}), \end{aligned} \tag{9.7}$$

making it right multiplication by an elliptic difference operator as required. □

As a result, there is a natural notion of leading coefficient for elements of $\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2)$.

Lemma 9.6. *There is an exact sequence*

$$0 \longrightarrow \hat{\mathcal{S}}'_{\rho;q;C}(0, D - C_m) \xrightarrow{T} \hat{\mathcal{S}}'_{\rho;q;C}(0, D) \xrightarrow{[T^0]} \Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))). \tag{9.8}$$

Proof. That the leading coefficient lies in the appropriate bundle follows by comparing leading coefficients when multiplying by operators of sufficiently (universally) ample degree. It thus remains only to show that if the leading coefficient vanishes, then we can factor out T. But this follows immediately from the corresponding statement for the original $\mathbb{Z}^{m+2}$-algebra. □

As one might expect, we do not in general need to test the entire translate of the universal nef cone.

Lemma 9.7. *Let D, D' be divisor classes such that D' is universally nef, with $D' \cdot C_m \ge 2$. If $\mathcal{D}$ is an elliptic difference operator such that*

$$\mathcal{S}'_{\rho;q;C}(D, D + D')\mathcal{D} \subset \mathcal{S}'_{\rho;q;C}(0, D + D'), \tag{9.9}$$

then $\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(0, D)$.

Proof. Let D'' be any universally very ample divisor class. Then

$$[D''+D']\mathcal{D} = [D''][D']\mathcal{D} \subset [D''][D+D'] \subset [D+D'+D''] \tag{9.10}$$

as required. □

Lemma 9.8. *For $m \geq 1$, if $D \cdot e_m = 0$, then*

$$\hat{\mathcal{S}}'_{\eta,x_0,\dots,x_m;q;C}(0,D) = \hat{\mathcal{S}}'_{\eta,x_0,\dots,x_{m-1};q;C}(0,D). \tag{9.11}$$

Proof. If $\mathcal{D} \in \hat{\mathcal{S}}'_{\eta,x_0,\dots,x_{m-1};q;C}(0,D)$, then we certainly have $[D']\mathcal{D} \subset [D+D']$ for some divisor D' which is universally very ample for X_{m-1}. But this divisor remains universally nef after blowing up, and thus $\mathcal{D} \in \hat{\mathcal{S}}'_{\eta,x_0,\dots,x_m;q;C}(0,D)$.

Conversely, if $\mathcal{D} \in \hat{\mathcal{S}}'_{\eta,x_0,\dots,x_m;q;C}(0,D)$, then there exists D' such that (a) $X_{m-1}(D') \cdot C_m \geq 2$, and (b) D' is universally nef with $[D']\mathcal{D} \subset [D+D']$, since (a) and (b) both contain translates of the universal nef cone. But then the claim follows by restricting ones attention to those operators of degree D' for which the first $D' \cdot e_m$ coefficients vanish. □

Remark. This is essentially the same argument that showed in the $m = 0$ cases that it was enough to saturate with respect to f.

Lemma 9.9. *Let D be a universally nef divisor class with $D \cdot C_m > 0$. Then*

$$\hat{\mathcal{S}}'_{\rho;q;C}(0,D) = \mathcal{S}'_{\rho;q;C}(0,D). \tag{9.12}$$

Proof. We first note that we may assume that either $m = 0$ (where the claim is true for *all* divisors) or $D \cdot e_m > 0$, since otherwise we may instead saturate in the blown down $\mathbb{Z}^{m+1}$-algebra.

If $D_2 - D_1$ is a multiple of f, then $\hat{\mathcal{S}}'_{\rho;q;C}(D_1,D_2)$ consists of operators of degree 0, and the corresponding functions must not only lie in the correct leading coefficient bundle, but must also have the correct symmetry; the claim follows immediately.

Otherwise, the leading coefficient map is already surjective before saturation, and thus we reduce to the claim for $D - C_m$, so eventually to smaller m. □

This lets us show that we get the same result if we saturate on the other side.

Corollary 9.10. *Let D be any divisor class. Then $\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(0,D)$ iff for all D' in some translate of the universal nef cone,*

$$\mathcal{D}\mathcal{S}'_{\rho;q;C}(-D',0) \subset \mathcal{S}'_{\rho;q;C}(-D',D). \tag{9.13}$$

Proof. Suppose first that $\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(0,D)$. Then it is certainly true that

$$\mathcal{D}\hat{\mathcal{S}}'_{\rho;q;C}(-D',0) \subset \hat{\mathcal{S}}'_{\rho;q;C}(-D',D) \tag{9.14}$$

for all D'. If $D+D'$ is universally nef with $(D+D')\cdot C_m > 0$, then $\hat{\mathcal{S}}'_{\rho;q;C}(-D',D) = \mathcal{S}'_{\rho;q;C}(-D',D)$ and thus

$$\mathcal{D}\mathcal{S}'_{\rho;q;C}(-D',0) \subset \mathcal{D}\hat{\mathcal{S}}'_{\rho;q;C}(-D',0) \subset \hat{\mathcal{S}}'_{\rho;q;C}(-D',D) = \mathcal{S}'_{\rho;q;C}(-D',D) \tag{9.15}$$

as required.

Taking into account the translation symmetry, we have shown that if an operator $\mathcal{D}$ satisfies

$$\mathcal{S}'_{\rho;q;C}(D_2,D_2+D')\mathcal{D} \subset \mathcal{S}'_{\rho;q;C}(D_1,D_2+D') \tag{9.16}$$

for all D' in some translate of the universal nef cone, then

$$\mathcal{D}\mathcal{S}'_{\rho;q;C}(D_1 - D', D_1) \subset \mathcal{S}'_{\rho;q;C}(D_1 - D', D_2) \tag{9.17}$$

for all D' in some translate of the universal nef cone. In particular, the converse is just the adjoint of what we have already shown! □

Remark. In particular, we see that the adjoint symmetry extends to the saturated $\mathbb{Z}^{m+2}$-algebra.

We also obtain the following slight improvement of Lemma 9.7. (The difference is that the product is only required to be in the saturated $\mathbb{Z}^{m+2}$-algebra, rather than the unsaturated $\mathbb{Z}^{m+2}$-algebra.)

Lemma 9.11. *Let D, D' be divisor classes such that D' is universally nef, with $D' \cdot C_m \geq 2$. If $\mathcal{D}$ is an elliptic difference operator such that*

$$\mathcal{S}'_{\rho;q;C}(D, D + D')\mathcal{D} \subset \hat{\mathcal{S}}'_{\rho;q;C}(0, D + D'), \tag{9.18}$$

then $\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(0, D)$.

Proof. Let D'' be a universally nef divisor class such that $D + D' + D''$ is universally nef with $(D + D' + D'') \cdot C_m > 0$. Then

$$[D'' + D']\mathcal{D} \subset [D'][D']\mathcal{D} \subset [D'']\hat{\mathcal{S}}'_{\rho;q;C}(0, D+D') \subset \hat{\mathcal{S}}'_{\rho;q;C}(0, D+D'+D'') = \mathcal{S}'_{\rho;q;C}(0, D+D'+D''). \tag{9.19}$$

Since $D' + D''$ is universally nef with $(D' + D'') \cdot C_m \geq 2$, the result follows from Lemma 9.7. □

Although the above is enough to give a reasonably sensible notion of sheaf, we still need to compare it to other notions of torsion. The most important of these is the following.

Definition 9.4. A homogeneous element v of a $\mathcal{S}'_{\eta,x_0,\dots,x_m;q;C}$-module is *inductively torsion* if there exists an integer r_m such that for any D with $D \cdot e_m \geq r_m$, every element of $[D]v$ is torsion as an element of the corresponding $\mathcal{S}'_{\eta,x_0,\dots,x_{m-1};q;C}$-module.

Here, we should really say "inductively torsion" in the condition on $[D]v$, but luckily this is an unnecessary distinction.

Proposition 9.12. *A homogeneous element v is inductively torsion iff it is torsion.*

Proof. Let v be inductively torsion, and let D_1 be any universally very ample divisor class such that $D_1 \cdot e_m \geq r_m$, so that $[D_1]v$ is torsion over $\mathcal{S}'_{\eta,x_0,\dots,x_{m-1};q;C}$. It follows that there is a divisor D_2 with $D_2 \cdot e_m = 0$ such that for any universally nef divisor D_3 with $D_3 \cdot e_m = 0$, $[D_2 + D_3][D_1]v = 0$. Again, we may choose D_2 so that Proposition 8.8 gives $[D_2 + D_3][D_1] = [D_1 + D_2 + D_3]$ for all D_3, and thus $[D_1 + D_2 + D_3]v = 0$, so that v is torsion.

Now, suppose $[D]v = 0$ for some universally very ample divisor class, and let D' be any divisor class with $D' \cdot e_m > D \cdot e_m$. It will suffice to show that every element of $[D']v$ is torsion over $\mathcal{S}'_{\eta,x_0,\dots,x_{m-1};q;C}$. Let D'' be a universally very ample divisor for X_{m-1}, and note that we may choose D'' so that $D' + D'' - D$ is universally very ample for X_m. But then

$$[D''][D']v \subset [D' + D'']v = [D' + D'' - D][D]v = 0, \tag{9.20}$$

so that every element of $[D']v$ is torsion as required. □

The significance of this is that "inductively torsion" is essentially just the usual notion of torsion for a sheaf $\mathbb{Z}$-algebra, in this case a sheaf $\mathbb{Z}$-algebra over $\mathcal{S}'_{\eta,x_0,\dots,x_{m-1};q;C}$.

Theorem 9.13. *The noncommutative scheme $X_{\eta,x_0,\dots,x_m;q;C}$ is the blowup in the sense of [64] of $X_{\eta,x_0,\dots,x_{m-1};q;C}$ at the point x_m. In particular, $X_{\eta,x_0,\dots,x_m;q;C}$ is a noetherian quasi-scheme.*

Proof. As we have already mentioned, $\operatorname{qcoh} X_{\eta,x_0,\dots,x_m;q;C}$ is the category of sheaves on a sheaf $\mathbb{Z}$-algebra over $X_{\eta,x_0,\dots,x_{m-1};q;C}$. Moreover, since every Hom space in $\mathcal{S}'_{\eta,x_0,\dots,x_m;q;C}$ is a subspace of a corresponding Hom space in $\mathcal{S}'_{\eta,x_0,\dots,x_{m-1};q;C}$, the Hom sheaves in this sheaf $\mathbb{Z}$-algebra are all 2-sided ideals. Moreover, this sheaf $\mathbb{Z}$-algebra is generated in degree 1 (i.e., $-e_m$): to see this, we need simply note that there is a divisor $D_0 = X_{m-1}(D_0)$ such that for any universally nef divisor $D = X_{m-1}(D)$,

$$[rD_0 + D - re_m] = [D_0 - e_m][(r-1)D_0 + D - (r-1)e_m], \tag{9.21}$$

and thus the ideal of degree r is the product of the ideal of degree $r-1$ and the ideal of degree 1. Each degree 1 ideal is cut out by the condition that the leading coefficient vanish at the appropriate point, and is thus the ideal sheaf of that point.

In other words, this sheaf $\mathbb{Z}$-algebra is precisely the sheaf $\mathbb{Z}$-algebra appearing in Van den Bergh's noncommutative blowup. (There is a technical point to be addressed below: we need to show that the point being blown up is contained in a commutative "divisor". The argument for this will use Corollary 9.26, which does not depend on $X_{\rho;q;C}$ being a blowup or even noetherian.) Since the blowup construction preserves the noetherian quasi-scheme property (and this is known for noncommutative Hirzebruch surfaces [14]), the remaining claim follows by induction. □

Remark 1. Strictly speaking, Van den Bergh's construction uses a graded sheaf-algebra rather than a sheaf $\mathbb{Z}$-algebra, but it is easy to see that the divisor condition lets one identify the $\mathbb{Z}$-algebra with one coming from a graded algebra.

Remark 2. The same argument shows that $X_{\eta,x_0,\dots,x_m;q;C}$ is *strongly* noetherian in the sense of [6, 14]: it remains noetherian under base change to any (commutative) noetherian k-algebra.

Corollary 9.14. *When $q = 1$, the quasi-scheme $X_{\eta,x_0,\dots,x_m;1;C}$ is a commutative rational surface (the iterated blowup of $F_1 \cong X_{\eta,x_0;1;C}$ in the points $x_1, \dots, x_m$).*

It thus makes sense to define "coherent" sheaves: these are precisely the noetherian objects of the category of quasicoherent sheaves.

Corollary 9.15. *For any divisor class D, the sheaf $\mathcal{O}_X(D)$ is coherent.*

Proof. That this holds for $\mathcal{O}_X$ follows from the fact that $X_{\rho;q;C}$ is a noetherian quasi-scheme. The translation symmetry induces an equivalence of sheaf categories taking the sheaf $\mathcal{O}_X$ to $\mathcal{O}_X(D)$, which must therefore also be noetherian. □

The Theorem tells us that the quasi-schemes we have constructed are not new; however, our construction gives a great deal more control over the category, and thus allows us to prove new symmetries.

Proposition 9.16. *If x_0/x_1 is not a power of q, then the quasi-scheme $X_{\eta,x_0,x_1;q;C}$ is isomorphic to the quasi-scheme $X_{x_0\eta/x_1,x_1,x_0;q;C}$, in such a way that the isomorphism preserves the anticanonical curve.*

Proof. Since the corresponding $\mathbb{Z}^3$-algebras are related by an isomorphism (essentially the Fourier transform), we simply need to show that the two notions of torsion are equivalent. In particular, the divisor $D_1 = 3s + 5f - 2e_1$ is invariant under that symmetry; thus if we can show that an element is torsion iff $[rD_1]v = 0$ for some $r > 0$, the claim will follow. Note that D_1 is universally nef with $D_1 \cdot C_1 = 11 \geq 2$, and thus if $[rD_1]v = 0$ for some $r > 0$, then v is torsion. It thus remains only to show the converse. Let $D_0 = 2s + 4f - e_1$. Since D_0 is universally very ample, any homogeneous torsion element will be annihilated by rD_0 for $r \gg 0$. If we can show that $[r(D_1 - D_0)][rD_0] = [rD_1]$, then the claim will follow. Since $[(r-1)D_1][D_1] = [rD_1]$ and $[D_1][(r-1)D_0] = [(r-1)D_0 + D_1] = [(r-1)D_0][D_1]$, we may immediately reduce to the case $r = 1$, which follows from the calculation

$$[s+f-e_1][2s+4f-e_1] = [s+f-e_1][f][2s+3f-e_1] = [s+2f-e_1][2s+3f-e_1] = [3s+5f-e_1]; \quad (9.22)$$

here the first and last steps are just Lemma 8.3, while the middle step is Lemma 8.16. □

Proposition 9.17. *If η/x_1x_2 is not a power of q, then the quasi-schemes $X_{\eta,x_0,x_1,x_2;q;C}$ and $X_{\eta,x_0\eta/x_1x_2,\eta/x_2,\eta/x_1;q;C}$ are isomorphic, in such a way that the isomorphism preserves the anti-canonical curve.*

Proof. The same argument applies, with $D_1 = 6s + 12f - 4e_1 - 2e_2$, $D_0 = 4s + 7f - 2e_1 - e_2$, and using Lemma 8.12 in place of Lemma 8.16. □

Proposition 9.18. *If x_{m-1}/x_m is not a power of q, then the quasi-schemes $X_{\eta,x_0,\dots,x_{m-2},x_{m-1},x_m;q;C}$ and $X_{\eta,x_0,\dots,x_{m-2},x_m,x_{m-1};q;C}$ are isomorphic, in such a way that the isomorphism preserves the anticanonical curve.*

Proof. Let D be a divisor which is in the interior of the facet of the universal nef cone cut out by $D \cdot (e_{m-1} - e_m)$; replacing D by $3D$ as necessary, we may assume $D \cdot C_m \geq 3$. Let $D' = 2s + af - e_1 - \cdots - e_m$ for $a \gg 0$. Then $D_1 = D + 2D' + e_{m-1} + e_m$ is still in the interior of the relevant facet, and $D_0 = D + D' + e_m$ is universally very ample. It thus suffices as before to show that $[D_1 - D_0][D_0] = [D_1]$, or in other words that

$$[D' + e_{m-1}][D + D' + e_m] = [D + 2D' + e_{m-1} + e_m]. \quad (9.23)$$

This follows from the calculation

$$[D'+e_{m-1}][D+D'+e_m] = [D'+e_{m-1}][D'+e_m][D] = [2D'+e_{m-1}+e_m][D] = [D+2D'+e_{m-1}+e_m]. \quad (9.24)$$

Here the first and third steps are Proposition 8.8, while the middle step is Lemma 8.11. □

Combining these gives us the following result.

Theorem 9.19. *Let r_i be any simple root of E_{m+1}, with corresponding reflection s_i. If $\rho(r_i)$ is not a power of q, then $X_{\rho;q;C} \cong X_{\rho\circ s_i;q;C}$.*

Proof. The only thing left to observe is that the symmetries we have shown all preserve the anti-canonical curve, and thus will continue to hold after blowing up. □

Call an element $w \in W(E_{m+1})$ *admissible* for ρ if there is some reduced word for w such that each simple reflection in that word induces an equivalence under the Theorem, starting with $X_{\rho;q;C}$, thus giving an isomorphism $X_{\rho;q;C} \cong X_{\rho\circ w^{-1};q;C}$.

There is an alternate description that does not depend on a choice of reduced word.

Proposition 9.20. *An element w is admissible iff no positive root that changes sign under w is taken to a power of q by ρ.*

Proof. The claim is certainly true if w is the identity, and by definition when w is a simple reflection. More generally, let $w = s_1 s_2 \cdots s_l$ be any reduced word for w. By induction on the length of w, we may assume that the claim is true for $s_1 w$; that is, that $s_1 w$ is admissible for precisely those ρ satisfying the given condition.

Suppose $s_1 w$ is admissible for ρ, or equivalently that no positive root that changes sign under $s_1 w$ maps to a power of q under ρ. By standard Coxeter theory, the set of positive roots made negative by w consists of those roots made negative by $s_1 w$ together with the root that w maps to the simple root corresponding to s_1. Thus s_1 is admissible for $\rho \circ (s_1 w)^{-1}$ iff ρ does not map that root to a power of q.

In other words, w is admissible for ρ relative to the given reduced expression iff it satisfies the desired invariant description; thus the latter is equivalent to the admissibility of w with respect to any reduced expression. □

Since the above symmetries do not preserve the universal nef cone (and thus do not preserve the set of universally very ample divisor classes), we obtain a large set of other notions of "torsion" by applying admissible elements.

Definition 9.5. A divisor class D is *nef* for $X_{\rho;q;C}$ if for some admissible element $w \in W(E_{m+1})$, wD is universally nef.

Note that if $q = 1$, this agrees with the usual notion of "nef" (having nonnegative intersection with every curve); we will show below that there is an analogous description in our setting. As this will also show that the nef divisors form a semigroup (the integer points of a suitable cone), we will refer to that set as the "nef cone". (This can also be shown directly from Coxeter theory.)

Lemma 9.21. *Let D be a nef divisor class with $D \cdot C_m \geq 2$. Then any homogeneous element v with $[D]v = 0$ is torsion.*

Proof. The notion of "torsion" is invariant under admissible simple reflections, and thus we may reduce to the case that D is in the universal nef cone. □

Ideally, we would have a result saying that we can replace the universal nef cone with the nef cone when defining "torsion". This has difficulties with the fact that the nef cone is in general a union of infinitely many images of the universal nef cone. To make things more manageable, it is useful to introduce an intermediate notion. Suppose S is a parabolic subsystem of E_{m+1} such that $W(S)$ is finite and every element of $W(S)$ is admissible. Then we say that D is "S-nef" if some point of its $W(S)$-orbit is universally nef; note that these elements form a cone. We say that D is "S-very ample" if it is contained in the interior of the S-nef cone and $D \cdot C_m \geq 3$.

Proposition 9.22. *Let S be as above, and let v be a homogeneous element of a $\mathcal{S}'_{\rho;q;C}$-module. Then the following are equivalent:*

(1) v is torsion.

(2) There is some translate of the S-nef cone annihilating v.

(3) There is some $W(S)$-invariant translate of the S-nef cone annihilating v.

(4) There is some S-very ample divisor D such that $[D]v = 0$.

(5) There is some $W(S)$-invariant S-very ample divisor D such that $[D]v = 0$.

Proof. That $3 \implies 2$ and $5 \implies 4$ are tautologies, that $2 \implies 4$ and $3 \implies 5$ follows from the fact that any two translates of the S-nef cone intersect (and thus any translate contains an S-very ample divisor), while $5 \implies 3$ by Proposition 8.8. Moreover, $4 \implies 1$ since any S-very ample divisor is nef. It thus remains only to show that $1 \implies 5$.

Suppose first that $s - e_1 \in S$ (so, in particular, $s - e_1$ is admissible). Let D be an $W(S)$-invariant S-very ample divisor D such that $D - 2f$ is S-very ample. Note that $D - f$ is then also S-very ample, and invariant under a smaller subgroup; there is thus an admissible element taking $D - f$ to the universal nef cone and giving an S'-very ample element for some $S' \subsetneq S$. It follows (by induction on $|S|$) that there exists $r > 0$ such that $[r(D - f)]v = 0$. To finish, it will suffice to show that $[f][D - f] = [D]$. Since $D + af$ is nef for $a \in \{-2, -1, 0\}$, we have $\mathcal{S}'_{\rho;q;C}(0, D + af) = \hat{\mathcal{S}}'_{\rho;q;C}(0, D + af)$. It follows in particular that as a varies, $\mathcal{S}'_{\rho;q;C}(0, D + af)$ is the saturated module representing a sheaf on $\mathbb{P}^1$. Since

$$\dim \mathcal{S}'_{\rho;q;C}(0, D) - 2 \dim \mathcal{S}'_{\rho;q;C}(0, D - f) + \dim \mathcal{S}'_{\rho;q;C}(0, D - 2f) = 0, \tag{9.25}$$

and thus the same holds for the saturated spaces, we immediately conclude that $[f][D - f] = [D]$.

We may thus reduce to the case that $s - e_1 \notin S$, so that S is contained in the root system D_m. Let w_0 be the longest element of $W(S)$, and for $1 \leq i \leq m$, let F_i denote the divisor $C_i + (m-1)f$. The divisor $D = \sum_{1 \leq i \leq m} (F_i + w_0 F_i)$ is $W(S)$-invariant, and we claim that $[rD]v = 0$ for $r \gg 0$. This is certainly true for the universally ample divisor $\sum_i F_i$, and thus for its image under w_0. It will thus suffice to show

$$[\sum_i w_0 F_i][\sum_i F_i] = [D]. \tag{9.26}$$

Since

$$[\sum_i F_i] \supset \prod_i [F_i] \tag{9.27}$$

and similarly for $[\sum_i w_0 F_i]$, it will suffice to show that we can pair up the factors so that the product for each pair is surjective and results in a universally nef degree (which by inspection has $X_5(D) \neq 0$). We may then apply Proposition 8.7 to move those factors to the end and induct. That $[F_i][w_0 F_i] = [F_i + w_0(F_i)]$ for each i follows from Corollary 8.15, Lemma 8.13, or Lemma 8.11, depending on how the component(s) (if any) of S not fixing F_i meet $\{f - e_1 - e_2, e_1 - e_2\}$. □

Call a divisor class "ample" (for $X_{\rho;q;C}$) if it is in the interior of the nef cone.

Lemma 9.23. *Any ample divisor class satisfies $D \cdot C_m > 0$ and $D \cdot e_m > 0$.*

Proof. It is equivalent to show that any nef divisor class satisfies $D \cdot C_m, D \cdot e_m \geq 0$. The first claim is trivial, since C_m is invariant under $W(E_{m+1})$. For the second claim, we need simply observe as in the proof of [47, Thm. 3.4] that for any element $w \in W(E_{m+1})$, $we_m - e_m$ is a nonnegative sum of roots. It follows that $D \cdot we_m \geq 0$ for any universally nef divisor class, and thus $w^{-1}D \cdot e_m \geq 0$; since w was arbitrary, this in particular applies to nef divisors. □

The name "ample" is justified by the following results.

Theorem 9.24. *Let D be an ample divisor class, and let v be a homogeneous element of a $\mathcal{S}'_{\rho;q;C}$-module. Then v is torsion iff $[rD]v$ for all sufficiently large r iff $[rD - \deg(v)]v = 0$ for all sufficiently large r.*

Proof. As usual, we may assume that D is universally nef. Let S be the parabolic subsystem of roots orthogonal to D. If we can show that $W(S)$ is finite and admissible, then any torsion element will be annihilated by some $W(S)$-invariant translate of the S-nef cone, and that cone will contain $D' + rD$ for all D' and all sufficiently large r (depending on D'). Conversely, rD and $rD - \deg(v)$ will both be S-very ample for sufficiently large r, giving the desired equivalence.

Since $D \cdot C_m, D \cdot e_m > 0$, the only facets of the universal nef cone that might contain D are those corresponding to simple reflections; in particular, the open face containing D is precisely the interior of the subcone of $W(S)$-invariant elements. Now, let A_D denote the set of directions in $\mathbb{R}\langle S\rangle$ for which infinitesimal deformations of D remain inside the nef cone. This contains the fundamental chamber of S, and all conjugates of the fundamental chamber by admissible elements. In particular, it is contained in the Tits cone of $W(S)$. Since D is an interior point, $A = \mathbb{R}\langle S\rangle$, and must in particular equal the Tits cone. It follows that $W(S)$ is finite (so that the Tits cone is the entire space), and that every element of $W(S)$ is admissible (since otherwise A would be missing some chambers). But this was precisely what we needed to show. □

Theorem 9.25. *Let D be an ample divisor class. Then for any sheaf M on $X_{\rho;q;C}$, the natural map*

$$\bigoplus_{r \geq R} \mathcal{O}_X(-rD) \otimes \mathrm{Hom}(\mathcal{O}_X(-rD), M) \to M \tag{9.28}$$

is surjective for all R. If M is coherent, then the term $r = R$ suffices for $R \gg 0$.

Proof. The category $\mathrm{qcoh}\, X_{\rho;q;C}$ is generated by the set of *all* line bundles (since $\mathcal{S}'_{\rho;q;C} - \mathcal{M}od$ is generated by the corresponding projective modules). As a result, we need only prove the Theorem in the case that $M = \mathcal{O}_X(D')$. In other words, we need to show that for each D', there exists r such that for all D'' in some translate of the universal nef cone, the map

$$\mathcal{S}'_{\rho;q;C}(rD, D'') \otimes \hat{\mathcal{S}}'_{\rho;q;C}(-D', rD) \to \mathcal{S}'_{\rho;q;C}(-D', D'') \tag{9.29}$$

is surjective. Moreover, by Theorem 9.24, it suffices to prove this for $D'' = sD$ for $s \gg 0$, where as usual we assume that D is universally nef. We need simply take r large enough that $rD + D'$ is nef (and can moreover arrange for it to be universally nef using admissible symmetries fixing D), with $(rD + D') \cdot C_m \geq 3$, at which point surjectivity holds whenever $s \geq r + 3$. □

The following is immediate. (See, e.g., [5].)

Corollary 9.26. *For any ample divisor class D, there is an isomorphism*

$$X_{\rho;q;C} \cong \mathrm{Proj}(\mathcal{S}'_{\rho;q;C}|_{\mathbb{Z}D}) \tag{9.30}$$

Remark. One special case of particular interest is the case $D = C_m$ for $m < 8$. This is only ample when all of $W(E_{m+1})$ is admissible, or in other words when $\rho(r) \notin q^{\mathbb{Z}}$ for all roots r. In this case, since each Hom space in the $\mathbb{Z}$-algebra has a canonical section (T^m), we may identify it with an actual graded algebra, which will be a flat deformation of the anticanonical embedding of a del Pezzo surface.

Corollary 9.27. *In the case $C = \mathbb{C}^*/\langle p\rangle$, the $\mathcal{S}'_{\rho;q;p}$-module $\mathcal{M}er$ is the module of global sections of a quasi-coherent sheaf.*

Proof. We need to prove that $\mathcal{M}er$ does not contain any torsion elements, and is saturated. By Corollary 9.26, it suffices to check this on the cosets of $\mathbb{Z}D$ for any ample divisor class D. Taking $D = (m(m+1)/2)f + \sum_{1 \leq k \leq m} C_k$, we find that for any $r > 0$, $[rD]$ contains operators of the form $g(z)T^{rm}$. The claim follows immediately from the fact that such operators act invertibly on the space of meromorphic functions. □

When q is torsion, we have already seen that the $\mathbb{Z}^{m+2}$-algebra $\mathcal{S}'_{\rho;q;C}$ contains a large "commutative" sub-$\mathbb{Z}^{m+2}$-algebra. We have, in fact, the following.

Proposition 9.28. *Let $q \in \mathrm{Pic}^0(C)$ be r-torsion, with corresponding isogeny $\phi : C \to C/\langle q\rangle =: C'$, and let Y denote the commutative rational surface corresponding to $X_{\phi_*\circ\rho;1;C'}$. Then there is a coherent, locally free $\mathcal{O}_Y$-algebra $\mathcal{A}$ such that $\operatorname{coh} X_{\rho;q;C}$ is naturally equivalent to the category of $\mathcal{A}$-modules, in such a way that $\mathcal{O}_X \mapsto \mathcal{A}$. Moreover, $\mathcal{A}$ is an Azumaya algebra (of degree r) on the complement of C'.*

Proof. Fix a divisor $D = af + \sum_{1\le i\le m} C_l$ with a sufficiently large that Lemma 8.17 applies. Such a divisor is certainly ample, and thus by Corollary 9.26, we have

$$X_{\rho;q;C} \cong \mathrm{Proj}(\mathcal{S}'_{\rho;q;C}|_{\mathbb{Z}rD}) \tag{9.31}$$

and similarly

$$Y \cong \mathrm{Proj}(\mathcal{S}'_{\phi_*\circ\rho;1;C'}|_{\mathbb{Z}D}). \tag{9.32}$$

Note that here, although we are working with $\mathbb{Z}$-algebras, we could equally well consider them as graded algebras; since $q^r = 1$, the translation by rD (resp. D) symmetry is an automorphism. Lemma 8.17 implies that the graded algebra corresponding to $\mathcal{S}'_{\rho;q;C}|_{\mathbb{Z}rD}$ is finitely generated over the graded algebra corresponding to $\mathcal{S}'_{\phi_*\circ\rho;1;C'}|_{\mathbb{Z}D}$, and thus gives rise to a coherent $\mathcal{O}_Y$-algebra as required.

Now, if we blow up a further point of $X_{\rho;q;C}$, thus also blowing up a point of Y, the resulting $\mathcal{O}_{\tilde{Y}}$-algebra is isomorphic to $\mathcal{A}$ on the complement of the exceptional locus. In particular, if we remove C' and the divisors $s, e_1, \dots, e_m$, the result will be the same as if we instead removed the anticanonical curve from the center of the appropriate noncommutative $\mathbb{P}^2$. It follows from [4, Thm. 7.3] that $\mathcal{A}$ is an Azumaya algebra on the complement of $C' \cup s \cup e_1 \cup \cdots \cup e_m$. This can be rephrased more canonically: For any birational morphism $\psi : Y \to \mathbb{P}^2$, $\mathcal{A}$ is an Azumaya algebra on $\psi^{-1}(\mathbb{P}^2 \setminus C')$. In fact, this extends to rational maps: if $\psi : Y \to \mathbb{P}^2$ is a rational map which is regular outside C', then $\mathcal{A}$ is an Azumaya algebra on $\psi^{-1}(\mathbb{P}^2 \setminus C')$. Indeed, such a rational map factors through an iterated blowup through points of C'. It is straightforward to verify that the open sets of the form $\psi^{-1}(\mathbb{P}^2 \setminus C')$ cover $Y \setminus C'$, and thus $\mathcal{A}$ is indeed an Azumaya algebra on $Y \setminus C'$.

It remains to show that $\mathcal{A}$ is locally free. This is immediate on $Y \setminus C'$, so it will suffice to show that the natural map $\mathcal{A}(-C') \to \mathcal{A}$ is injective and its cokernel is locally free as a sheaf on C'. (It then follows that $\mathrm{Tor}_d(\mathcal{A}, O_p) = 0$ for any $d > 0$ and any point $p \in C'$). The sheaf $\mathcal{A}(-C')$ corresponds to the graded module which in degree d is given by $\mathcal{S}'_{\rho;q;C}(0, drD - rC_m)$. It follows from the proof of Lemma 8.17 that this module is generated in degree 1, and that both it and $\mathcal{S}'_{\rho;q;C}(0, drD)$ are saturated over the homogeneous coordinate ring of Y. Injectivity thus follows immediately from injectivity for the graded modules, while the cokernel corresponds to the graded quotient module. This quotient module has a natural filtration with subquotients $[T^0]\mathcal{S}'_{\rho;q;C}(0, drD - lC_m)$ for $0 \le l < r$; since those subquotients are invertible sheaves on C and $\phi : C \to C'$ is finite and flat, we conclude that the subquotients are locally free sheaves on C', and the claim follows. □

Remark. We will generalize this in [44], and in particular prove that $\mathcal{A}$ is a maximal order in its generic fiber.

In the proof of Theorem 9.13, we postponed a technical point. Van den Bergh's construction only defines a blowup in a point of a commutative curve that appears as a divisor in a suitable sense; in particular, there must be an endofunctor corresponding to twisting by the curve.

In our case, the curve should clearly be C, and the corresponding 2-sided ideal in $\mathcal{S}'_{\rho;q;C}$ is the kernel of $[T^0]$, or equivalently the image of T. Since T is invertible as a difference operator, we may use it to define an endofunctor of $\mathcal{S}'_{\rho;q;C}$ as follows.

Definition 9.6. The "canonical functor" of $\mathcal{S}'_{\rho;q;C}$ is the endofunctor $\bar{\theta}$ which acts on objects as $D \mapsto D + C_m$ and on morphisms as $\mathcal{D} \mapsto T\mathcal{D}T^{-1}$.

Remark. We will introduce a slightly modified version of the same functor below.

This induces an endofunctor of $\mathcal{S}'_{\rho;q;C} - \mathcal{M}od$ (the categorical analogue of pulling back a module through a ring homomorphism), and since that endofunctor preserves the subcategory of torsion modules, it also induces an endofunctor of $\operatorname{qcoh} X_{\rho;q;C}$, which we also denote by $\bar{\theta}$. Note that $\bar{\theta}\mathcal{O}_X(D) \cong \mathcal{O}_X(D - C_m)$, and there is a natural transformation $T : \bar{\theta} \to \mathrm{id}$ which on line bundles is the operator

$$T \in \hat{\mathcal{S}}'_{\rho;q;C}(-D, -D + C_m) = \operatorname{Hom}(\mathcal{O}_X(D - C_m), \mathcal{O}_X(D)). \tag{9.33}$$

Note that in the commutative setting, $\bar{\theta}M \cong M \otimes \omega_X$, and T corresponds to a Poisson structure $T : \omega_X \to \mathcal{O}_X$. It turns out that T is nearly always unique.

Proposition 9.29. *If $q \neq 1$, then any natural transformation $\bar{\theta} \to \mathrm{id}$ is proportional to T.*

Proof. Such a natural transformation T' is determined by its action on line bundles, and thus on the corresponding operators

$$T'_D \in \hat{\mathcal{S}}'_{\rho;q;C}(D, D + C_m). \tag{9.34}$$

For these to give a natural transformation, they must satisfy

$$T^{-1}T'_{D_2}\mathcal{D} = \mathcal{D}T^{-1}T'_{D_1} \tag{9.35}$$

for any operator $\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2)$. When $D_2 - D_1 \in f$, so $\mathcal{D}$ is a multiplication operator, the fact that this holds for any $\mathcal{D}$ of that degree easily implies (since $q \neq 1$) that the operator $T^{-1}T'_{D_1}$ must itself be a multiplication operator. But then degree considerations imply that $T^{-1}T'_D$ is a scalar for all D. Since then $T^{-1}T'_{D_1} = T^{-1}T'_{D_2}$ whenever $\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2) \neq 0$, we see that $T' = \alpha T$ for some fixed scalar α as required. □

Now, the category $[T^0]\mathcal{S}'_{\rho;q;C}$ can also be viewed as a right $\mathcal{S}'_{\rho;q;C}$-module, and this gives rise to a pair of adjoint functors $[T^0]^*$ and $[T^0]_*$ (tensoring and taking Homs respectively). Equivalently, we may fix an ample divisor D and restrict both categories to $\mathbb{Z}D$. Since $[T^0]\mathcal{S}'_{\rho;q;C}|_{\mathbb{Z}D}$ is a twisted homogeneous coordinate ring of C, taking sheaves gives a pair of adjoint functors

$$i_* : \operatorname{qcoh} C \to \operatorname{qcoh} X_{\rho;q;C} \qquad \text{and} \qquad i^* : \operatorname{qcoh} X_{\rho;q;C} \to \operatorname{qcoh} C. \tag{9.36}$$

(It suffices to verify that $[T^0]_*$ and $[T^0]^*$ take torsion modules to torsion modules, and this is straightforward.) In addition, the leading coefficient filtration establishes $\mathcal{S}'_{\rho;q;C}|_{\mathbb{Z}D}$ as the global section module of a locally free sheaf on $[T^0]\mathcal{S}'_{\rho;q;C}|_{\mathbb{Z}D}$, and thus i_* is exact.

Proposition 9.30. *The functor i_* embeds C as a divisor in $X_{\rho;q;C}$ in the sense of [64, §3.7].*

Proof. First note that if a sheaf M is in the image of i_*, then the morphism $T_M : \bar{\theta}M \to M$ is 0; this simply follows from the fact that the operator T is in the kernel of $[T^0]$.

We next observe that $\mathcal{O}_X(-rD)/T\bar{\theta}\mathcal{O}_X(-rD)$ is in the image of i_* (as the corresponding quotient of projective modules is certainly in the image of $[T^0]_*$). It follows that if M is a sheaf such

that $T_M : \bar\theta M \to M$ is 0, then M is a quotient of a sheaf i_*N. Indeed, M is a quotient of a sum of sheaves $\mathcal{O}_X(-rD)$, and the naturality of T implies that any map $\mathcal{O}_X(-rD) \to M$ kills $T\bar\theta\mathcal{O}_X(-rD)$, so factors through a map $\mathcal{O}_X(-rD)/T\bar\theta\mathcal{O}_X(-rD)$.

Now, the property $T_M = 0$ is preserved under taking sub- and quotient sheaves. Thus the kernel of a surjection $i_*N \to M$ is itself a quotient of a sheaf in the image of i_*. Any morphism $i_*N' \to i_*N$ is itself in the image of i_*, and thus M is the image of the appropriate cokernel.

We thus see that M is in the image of i_* iff $T_M = 0$, and this establishes that $\bar\theta = o_X(-C)$ in the notation of [64], making C a divisor in $X_{\rho;q;C}$. □

Remark. In particular, we may indeed use [64] to blow up $X_{\rho;q;C}$ in any point of C. The identification of the relevant sheaf $\mathbb{Z}$-algebra with a graded algebra uses the functor $\bar\theta$, which establishes an isomorphism between the $\mathbb{Z}$-algebra and its shift.

In the sequel, we will simply silently apply i_*, identifying $\operatorname{qcoh} C$ with the corresponding subcategory of $\operatorname{qcoh} X_{\rho;q;C}$. The functor i^* will be denoted by $M|_C := i^*M$, where we note that

$$i_*i^*M \cong M/T\bar\theta M. \tag{9.37}$$

With this notation, we have the following as an immediate consequence of the adjunction.

Proposition 9.31. *For any object $N \in D^b \operatorname{coh} C$, $R\operatorname{Hom}(M, N) \cong R\operatorname{Hom}_C(M|_C^{\mathbf{L}}, N)$.*

Definition 9.7. A sheaf M on $X_{\rho;q;C}$ is *transverse to C* if $M|_C^{\mathbf{L}}$ is a sheaf, *disjoint from C* if $M|_C = 0$.

Proposition 9.32. *If M is supported on C and N is disjoint from C, then $R\operatorname{Hom}(M, N) = 0$.*

Proof. The commutative diagram

$$\begin{array}{ccc} R\operatorname{Hom}(M, \bar\theta N) & \xrightarrow{0} & R\operatorname{Hom}(\bar\theta M, \bar\theta N) \\ {\scriptstyle T_N}\downarrow & & {\scriptstyle T_N}\downarrow \\ R\operatorname{Hom}(M, N) & \xrightarrow{0} & R\operatorname{Hom}(\bar\theta M, N) \end{array} \tag{9.38}$$

remains commutative if we include the isomorphism $\bar\theta : R\operatorname{Hom}(M, N) \cong R\operatorname{Hom}(\bar\theta M, \bar\theta N)$. Since T_N is an isomorphism, this implies $R\operatorname{Hom}(M, N) = 0$. □

We now return to the saturated category, and in particular the question of computing the Hom spaces in that category (or, at the very least, their dimensions); we will see below that this algorithm easily refines to one for computing dimensions of Ext spaces between line bundles on $X_{\rho;q;C}$.

By analogy with the commutative case [47], our strategy will be to reduce to the case that the degree is in the fundamental chamber. We have already mostly dealt with this case: we saw above that when D is universally nef with $D \cdot C_m > 0$, then

$$\hat{\mathcal{S}}'_{\rho;q;C}(0, D) = \mathcal{S}'_{\rho;q;C}(0, D), \tag{9.39}$$

and thus flatness of the latter gives

$$\dim \hat{\mathcal{S}}'_{\rho;q;C}(0, D) = 1 + \frac{D \cdot (D + C_m)}{2} \tag{9.40}$$

by reduction to the commutative case.

The other possibilities in the fundamental chamber are that we could have $D \cdot e_m < 0$ or $D \cdot C_m \leq 0$.

Lemma 9.33. *If $D \cdot e_m = -d$ for $d > 0$, then $\hat{\mathcal{S}}'_{\rho;q;C}(0, D) = \hat{\mathcal{S}}'_{\rho;q;C}(0, D - de_m)$.*

Proof. Since $\hat{\mathcal{S}}'_{\rho;q;C}(0, D - de_m)$ can be computed on the blown down surface, we have

$$T^d \hat{\mathcal{S}}'_{\rho;q;C}(0, D - de_m) \subset \hat{\mathcal{S}}'_{\rho;q;C}(0, D - de_m + dC_{m-1}) = \hat{\mathcal{S}}'_{\rho;q;C}(0, D + dC_m) \tag{9.41}$$

and is thus in particular equal to the subspace of $\hat{\mathcal{S}}'_{\rho;q;C}(0, D + dC_m)$ where the first d coefficients vanish. But this is the same as $T^d \hat{\mathcal{S}}'_{\rho;q;C}(0, D)$, so the claim follows after dividing by T^d. □

Lemma 9.34. *If $D \cdot C_m < 0$, then $\hat{\mathcal{S}}'_{\rho;q;C}(0, D) = \hat{\mathcal{S}}'_{\rho;q;C}(0, D - C_m)$.*

Proof. This follows immediately from Lemma 9.6; the natural inclusion

$$\hat{\mathcal{S}}'_{\rho;q;C}(0, D - C_m) \to \hat{\mathcal{S}}'_{\rho;q;C}(0, D) \tag{9.42}$$

has quotient contained in

$$\Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))) = 0. \tag{9.43}$$

□

We are thus left with the case of a universally nef divisor class with $D \cdot C_m = 0$. For this, we have the following (surprisingly tricky) result.

Lemma 9.35. *Let D be a universally nef divisor class with $D \cdot C_m = 0$. Then there is a short exact sequence*

$$0 \longrightarrow \hat{\mathcal{S}}'_{\rho;q;C}(0, D - C_m) \xrightarrow{T} \hat{\mathcal{S}}'_{\rho;q;C}(0, D) \xrightarrow{[T^0]} \Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))) \longrightarrow 0. \tag{9.44}$$

Proof. We may as well assume $D \cdot e_m > 0$, as if not we can reduce to the blown down surface. Once $m < 8$, the only universally nef divisor with $D \cdot C_m = 0$ is $D = 0$, where the result is trivial.

Suppose now that $m > 8$. If we multiply an element of $\hat{\mathcal{S}}'_{\rho;q;C}(0, D)$ by any element of $\mathcal{S}'_{\rho;q;C}(-f, 0)$, the result is an element of $\hat{\mathcal{S}}'_{\rho;q;C}(-f, D) = \mathcal{S}'_{\rho;q;C}(-f, D)$. It follows that $\hat{\mathcal{S}}'_{\rho;q;C}(0, D)$ is the space of global sections of a suitable vector bundle on $\mathbb{P}^1$. Moreover, there is an induced leading coefficient exact sequence for that vector bundle, and the subbundle, corresponding to $\hat{\mathcal{S}}'_{\rho;q;C}(0, D - C_m) = \mathcal{S}'_{\rho;q;C}(0, D - C_m)$, is acyclic. (This holds because $(D - C_m) \cdot C_m > D \cdot C_m = 0$.) Since the quotient is just the direct image of $\Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))$, the claim follows.

We are thus left with the case $m = 8$, when $D = dC_8$ for some d. Now, we have an isomorphism

$$\mathcal{L}(\mathfrak{D}_{\rho;q;C}(dC_8)) \cong \mathcal{L}(\mathfrak{D}_{\rho;q;C}(C_8))^d, \tag{9.45}$$

so there are two natural cases to consider. If $\mathcal{L}(\mathfrak{D}_{\rho;q;C}(C_8))$ is non-torsion, then

$$\Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(dC_8))) = 0 \tag{9.46}$$

for all $d > 0$, and surjectivity is automatic. Otherwise, suppose it is r-torsion. In this case, we still trivially have surjectivity whenever d is not a multiple of r. When d is a multiple of r, we have $h^0(\mathcal{L}(\mathfrak{D}_{\rho;q;C}(dC_8))) = 1$, and thus the objective is simply to prove that there is an element of $\hat{\mathcal{S}}'_{\rho;q;C}(0, dC_8)$ with nonzero leading coefficient.

Clearly, it suffices to prove this in the case $d = r$. In that case, we find in the same way as for $m > 8$ that $\hat{\mathcal{S}}'_{\rho;q;C}(0, rC_8)$ is the space of global sections of a suitable vector bundle on $\mathbb{P}^2$. The corresponding filtration by leading degree has the sequence of subquotients

$$\pi_*(\mathcal{L}(\mathfrak{D}_{\rho;q;C}(rC_8))), \ldots \pi_*(\mathcal{L}(\mathfrak{D}_{\rho;q;C}(C_8))), \mathcal{O}_{\mathbb{P}^1} \tag{9.47}$$

where $\pi : C \to \mathbb{P}^1$ is the appropriate degree 2 morphism. Now, $\mathcal{O}_{\mathbb{P}^1}$ is certainly acyclic, and $\pi_*(\mathcal{L}(\mathfrak{D}_{\rho;q;C}(lC_8)))$ is acyclic whenever l is not a multiple of r. It follows, therefore, that the subbundle in the leading coefficient short exact sequence is acyclic as required. □

Corollary 9.36. *If the point $\eta^3 x_0^2/x_1\cdots x_8 \in \mathrm{Pic}^0(C)$ is r-torsion, then there is a difference operator $\mathcal{D}$ such that for any $d \geq 0$,*

$$T^{-rd}\hat{\mathcal{S}}'_{\eta,x_0,\ldots,x_8;q;C}(0, drC_8) = \langle 1, \mathcal{D}, \ldots, \mathcal{D}^d\rangle. \tag{9.48}$$

Proof. This is certainly true for $d = 1$ (bearing in mind that $\mathcal{D}$ will include negative shifts as well as positive shifts). It remains only to note that

$$T^{rd}\mathcal{D}T^{-r(d-1)} \in T^{r(d-1)}\mathcal{S}'_{\eta,x_0,\ldots,x_8;q;C}(0, rC_8)T^{-r(d-1)} = \mathcal{S}'_{\eta,x_0,\ldots,x_8;q;C}((d-1)rC_8, drC_8) \tag{9.49}$$

□

Remark. For $r = 1$, the operator $\mathcal{D}$ is the univariate case of an integrable Hamiltonian introduced by van Diejen [16] (see [43] for a discussion of this fact in terms of the Sklyanin algebra). It will be shown in [46] that this result generalizes to a multivariate statement, in which all of the relevant commuting operators appear. Another consequence is that the $W(E_9)$ symmetry induces a $W(E_8)$ symmetry on the parameters, which appears in the eigenvalue analysis of [53]. For $r > 1$, we obtain a new family of eigenvalue equations (now order $2r$), now with a symmetry of the form $W(E_8) \ltimes \Lambda_{E_8}/r\Lambda_{E_8}$, though it is unclear if this symmetry extends to the analytic setting of [53].

At this point, we can deal with any saturated Hom space with degree in the fundamental chamber for $W(E_{m+1})$. Thus, let α be a simple root of $W(E_{m+1})$; we wish to show that if $D\cdot\alpha < 0$, then $\hat{\mathcal{S}}'_{\rho;q;C}(0, D)$ is isomorphic to a space of the form $\hat{\mathcal{S}}'_{\rho';q;C}(0, D')$ with $D'\cdot\alpha > D\cdot\alpha$. This will clearly eventually result in a divisor with $D\cdot\alpha > 0$; more generally, given any set of simple roots that generate a finite root system, repeated reductions of this form will eventually result in a divisor in the fundamental chamber for that root system. In particular, if we do this for the subsystem D_m of roots orthogonal to f, none of the resulting reductions will change the degree of the corresponding operators, while the reduction corresponding to $s - e_1$ will always decrease that degree. Thus if we alternate reducing w.r.to D_m and reducing w.r.to $s - e_1$, every other step will strictly decrease $D \cdot f$. Thus such a process will either produce a divisor with $D \cdot f < 0$ (for which the saturated Hom space is necessarily empty) or one in the fundamental chamber.

The first case we need to deal with is when D is itself a multiple of α.

Lemma 9.37. *If $\rho(\alpha)$ is not a power of q, then $\hat{\mathcal{S}}'_{\rho;q;C}(0, r\alpha) = 0$ for all $r \in \mathbb{Z}$. Otherwise, if l is the smallest positive integer such that $\rho(\alpha) = q^l$, then $\dim \hat{\mathcal{S}}'_{\rho;q;C}(0, l\alpha) = 1$.*

Proof. There are three cases to consider: $\alpha = s - e_1$ for $m = 1$, $\alpha = f - e_1 - e_2$ for $m = 2$, or $\alpha = e_{m-1} - e_m$ for $m \geq 2$. In fact, the first two cases can be reduced to the third: for $s - e_1$, we need simply work on $\hat{\mathcal{S}}'_{\eta,x_0,x_1,x_2}$ for any x_2 for which x_0/x_2 is not a power of q (and thus x_1/x_2 is also not a power of q), and use the isomorphisms

$$\hat{\mathcal{S}}'_{\eta,x_0,x_1,x_2;q;C}(0, r(s - e_1)) \cong \hat{\mathcal{S}}'_{\eta,x_0,x_2,x_1;q;C}(0, r(s - e_2)) \cong \hat{\mathcal{S}}'_{x_0\eta/x_2,x_2,x_0,x_1;q;C}(0, r(e_1 - e_2)). \tag{9.50}$$

Similarly, for x_3 such that x_3/x_2 is not a power of q, we have

$$\begin{aligned}\hat{\mathcal{S}}'_{\eta,x_0,x_1,x_2,x_3;q;C}(0, r(f - e_1 - e_2)) &\cong \hat{\mathcal{S}}'_{\eta,x_0,x_1,x_3,x_2;q;C}(0, r(f - e_1 - e_3)) \\ &\cong \hat{\mathcal{S}}'_{\eta,x_0\eta/x_1x_3,\eta/x_3,\eta/x_1,x_2;q;C}(0, r(e_2 - e_3)).\end{aligned} \tag{9.51}$$

Now, we have

$$\hat{\mathcal{S}}'_{\rho;q;C}(0, r(e_{m-1} - e_m)) \subset \hat{\mathcal{S}}'_{\rho;q;C}(0, re_{m-1}) = \langle 1 \rangle, \tag{9.52}$$

and thus the only question is when $1 \in \hat{\mathcal{S}}'_{\rho;q;C}(0, r(e_{m-1} - e_m))$. It follows from the leading coefficient exact sequence that

$$\hat{\mathcal{S}}'_{\rho;q;C}(0, r(e_{m-1} - e_m)) \subset \Gamma(C; \mathcal{L}(\mathfrak{D}_{\rho;q;C}(r(e_{m-1} - e_m)))) \tag{9.53}$$

and thus a necessary condition is that

$$\mathfrak{D}_{\rho;q;C}(r(e_{m-1} - e_m)) = 0, \tag{9.54}$$

or in other words

$$\prod_{0 \le k < r} [q^{-k} x_{m-1}] = \prod_{0 \le k < r} [q^{r-k} x_m]. \tag{9.55}$$

It follows in particular that $x_{m-1} = q^{r-k} x_m$ for some $0 \le k < r$. (In fact, one can show that either $x_{m-1} = q^r x_m$ or $q^r = 1$.)

It remains only to show that when $x_{m-1} = q^l x_m$, with $l > 0$ the minimum solution, then $\mathcal{S}'_{\rho;q;C}(l(e_{m-1} - e_m), D) \subset \mathcal{S}'_{\rho;q;C}(0, D)$ for all D in some translate of the universal nef cone. It will be convenient to use the translation symmetry to reduce to showing that if $x_{m-1} = x_m$, then for all $l \ge 0$, $\mathcal{S}'_{\rho;q;C}(le_{m-1}, D) \subset \mathcal{S}'_{\rho;q;C}(le_m, D)$.

If $D \cdot (e_{m-1} - e_m) = 0$, then this follows immediately from the Weyl group symmetry. Now, choose such a divisor subject to the additional condition that $D - le_{m-1}$ is universally nef with $(D - le_{m-1}) \cdot C_m \ge 3$. Then for any universally nef D' with $D' \cdot C_m \ge 3$, we have the following:

$$\mathcal{S}'_{\rho;q;C}(le_{m-1}, D + D') = [D']\mathcal{S}'_{\rho;q;C}(le_{m-1}, D) = [D']\mathcal{S}'_{\rho;q;C}(le_m, D) \subset \mathcal{S}'_{\rho;q;C}(le_m, D + D'). \tag{9.56}$$

The claim follows. □

Remark. Of course, the case $s - e_1$ is just an elementary transformation away from the case $s - f$ we needed to consider in the F_0 case. In addition, if we unpack the symmetries we used, we find that $\hat{\mathcal{S}}'_{\rho;q;C}(0, l(s - e_1)) \subset \hat{\mathcal{S}}'_{\rho;q;C}(0, ls)$ and $\hat{\mathcal{S}}'_{\rho;q;C}(0, l(f - e_1 - e_2)) \subset \hat{\mathcal{S}}'_{\rho;q;C}(0, l(f - e_1))$, allowing us to identify the nonzero sections when they exist.

Lemma 9.38. *If $q^r = 1$ and α is a simple root with $\rho(\alpha)$ a power of q, then* $\dim \mathcal{S}'_{\rho;q;C}(0, r\alpha) = 1$.

Proof. We can either reduce to the commutative sub-$\mathbb{Z}^{m+2}$-algebra, or observe that both factors in

$$\hat{\mathcal{S}}'_{\rho;q;C}(0, l\alpha)\hat{\mathcal{S}}'_{\rho;q;C}(l\alpha, r\alpha) \tag{9.57}$$

are nonzero, where l is the minimum positive integer such that $\rho(\alpha) = q^l$. □

By analogy with the F_0 case, we expect that any space $\hat{\mathcal{S}}'_{\rho;q;C}(0, D)$ will either equal its image under the reflection, or will factor through one of the morphisms we have just constructed. It will be convenient as in the above proof to use the translation symmetry to arrange that $x_{m-1} = x_m$, so that ρ is invariant under the corresponding reflection.

Lemma 9.39. *Let $x_{m-1} = x_m$, and let D be any divisor with $D = X_{m-2}(D)$. Then for $0 \le l \le |\langle q \rangle|$, we have*

$$\hat{\mathcal{S}}'_{\rho;q;C}(0, D - le_m) = \hat{\mathcal{S}}'_{\rho;q;C}(0, D - le_{m-1}). \tag{9.58}$$

Proof. Using Lemma 9.11, we may as well assume that D is very far inside the universal nef cone of X_{m-2}, so that both spaces are subspaces of $\hat{\mathcal{S}}'_{\rho;q;C}(0,D) = \mathcal{S}'_{\rho;q;C}(0,D)$. Also, since $\hat{\mathcal{S}}'_{\rho;q;C}(D - le_{m-1}, D - le_m) = \langle 1\rangle$, what we need to establish is that every $\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(0, D - le_m)$ factors through this morphism.

Now, if the first l coefficients of $\mathcal{D}$ all vanish, then we find that

$$T^{-l}\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(0,(D-lC_{m-2})+le_{m-1}) = \hat{\mathcal{S}}'_{\rho;q;C}(0,D-lC_{m-2}) = \hat{\mathcal{S}}'_{\rho;q;C}(0,D-lC_{m-2}+le_m), \quad (9.59)$$

and thus the result follows. Thus, in order to prove the desired result, all we need to show is that for each $0 \le k < l$, the inclusion map

$$\hat{\mathcal{S}}'_{\rho;q;C}(0,D - le_{m-1} - kC_{m-1}) = \hat{\mathcal{S}}'_{\rho;q;C}(0,D - le_{m-1} - kC_m) \to \hat{\mathcal{S}}'_{\rho;q;C}(0,D - le_m - kC_m) \quad (9.60)$$

is surjective on leading coefficients. Since both spaces are contained in $\mathcal{S}'_{\rho;q;C}(0,D-kC_{m-2})$, it will suffice to show that

$$\mathfrak{D}_{\rho;q;C}(D - le_{m-1} - kC_{m-1})^{-1} = \gcd(\mathfrak{D}_{\rho;q;C}(D - le_m - kC_m)^{-1}, \mathfrak{D}_{\rho;q;C}(D - kC_{m-2})^{-1}). \quad (9.61)$$

Equivalently, we may multiply all three divisors by $\mathfrak{D}_{\rho;q;C}(D - kC_{m-2})$, so that we are comparing

$$\mathfrak{D}_{\rho;q;C}((k-l)e_{m-1})^{-1} \quad \text{and} \quad \gcd(\mathfrak{D}_{\rho;q;C}(ke_{m-1} + (k-l)e_m)^{-1}, 1). \quad (9.62)$$

Further multiplying by $\mathfrak{D}_{\rho;q;C}((k-l)e_{m-1})$ reduces to showing

$$\gcd(\prod_{-k\le j<0}[q^{-j}x_m], \prod_{0\le j<l-k}[q^{-j}x_m]) = 1. \quad (9.63)$$

But this is clearly true iff $l \le |\langle q\rangle|$. □

Lemma 9.40. *Suppose $x_{m-1} = x_m$, and let D_1, D_2 be divisors such that $|\langle q\rangle| \ge D_1\cdot(e_{m-1}-e_m) \ge 0 \ge D_2\cdot(e_{m-1}-e_m)$. Then*

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1,D_2) = \hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1,D_2) \quad (9.64)$$

Proof. Note that $\hat{\mathcal{S}}'_{\rho;q;C}(D_1, s_{m-1}D_1) = \langle 1\rangle$, we need only show that

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1,D_2) \subset \hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1,D_2); \quad (9.65)$$

i.e., that every morphism from D_1 to D_2 factors through $s_{m-1}D_1$.

If D' is universally nef with $D'\cdot C_m \ge 2$ and

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1,D_2+D') = \hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1,D_2+D'), \quad (9.66)$$

then we have

$$[D']\hat{\mathcal{S}}'_{\rho;q;C}(D_1,D_2) \subset \hat{\mathcal{S}}'_{\rho;q;C}(D_1,D_2+D') = \hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1,D_2+D'), \quad (9.67)$$

and thus Lemma 9.11 gives the desired result. In particular, we may freely add an arbitrary universally nef divisor with $D'\cdot(e_{m-1}-e_m) = -D_2\cdot(e_{m-1}-e_m)$ and $D'\cdot C_m \ge 2$ to D_2 before proving our claim. In particular, it suffices to prove the claim in the case $D_2\cdot(e_{m-1}-e_m) = 0$, $(D_2 - D_1)\cdot e_{m-1} \ge 0$. In addition, if $(D_2 - D_1)\cdot e_{m-1} > 0$, then we see that (adding a symmetric universally nef divisor class as necessary) both saturated Hom spaces surject onto their leading

coefficient bundles, and thus we may reduce to the case that the leading coefficient vanishes. In other words, it suffices to prove the result in the case $D_2 \cdot (e_{m-1} - e_m) = 0$, $(D_2 - D_1) \cdot e_{m-1} = 0$.

Using the translation freedom, we thus reduce to showing

$$\hat{\mathcal{S}}_{\rho;q;C}(le_m, D) \subset \hat{\mathcal{S}}_{\rho;q;C}(le_{m-1}, D) = \mathcal{S}_{\rho;q;C}(le_{m-1}, D) = \mathcal{S}_{\rho;q;C}(le_m, D), \tag{9.68}$$

where $D = X_{m-2}(D)$ is sufficiently far in the interior of the universal nef cone for $m-2$. But this is just the adjoint of the previous lemma. □

The adjoint is also useful.

Corollary 9.41. *Suppose $x_{m-1} = x_m$, and let D_1, D_2 be divisors such that $D_1 \cdot (e_{m-1} - e_m) \geq 0 \geq D_2 \cdot (e_{m-1} - e_m) \geq -|\langle q \rangle|$. Then*

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2) = \hat{\mathcal{S}}'_{\rho;q;C}(D_1, s_{m-1}D_2) \tag{9.69}$$

Lemma 9.42. *Suppose $x_{m-1} = x_m$, and let D_1, D_2 be divisors such that $-|\langle q \rangle| \leq D_2 \cdot (e_{m-1} - e_m), D_1 \cdot (e_{m-1} - e_m) \leq 0$. Then*

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2) = \hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1, s_{m-1}D_2). \tag{9.70}$$

Proof. The previous two results give

$$\hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1, D_2) = \hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2) \tag{9.71}$$

and

$$\hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1, D_2) = \hat{\mathcal{S}}'_{\rho;q;C}(s_{m-1}D_1, s_{m-1}D_2) \tag{9.72}$$

respectively. □

Similar results (using the above reduction) apply to the other simple roots, and we find the following (again assuming for simplicity that $\rho(\alpha) = 1$) for $\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2)$ when $(D_2 - D_1) \cdot \alpha < 0$:

(1) If $D_2 \cdot \alpha$ and $D_1 \cdot \alpha$ have opposite signs, or if there is a multiple of $|\langle q \rangle|$ between them, then

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2) = \hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2 - l\alpha), \tag{9.73}$$

where $l = (-D_2 \cdot \alpha) \bmod r$.

(2) Otherwise,

$$\hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2) = \hat{\mathcal{S}}'_{\rho;q;C}(s_\alpha D_1, s_\alpha D_2). \tag{9.74}$$

Here (2) is Lemma 9.42, while (1) is Corollary 9.41, translating by $|\langle q \rangle| e_m$ as necessary to ensure that $-|\langle q \rangle| \leq D_2 \cdot \alpha < 0$. Since (1) always increases $(D_2 - D_1) \cdot \alpha$ and (2) automatically terminates, iterating this process will always end with a pair D_1, D_2 such that $(D_2 - D_1) \cdot \alpha \geq 0$.

This completes the desired algorithm for computing $\dim \hat{\mathcal{S}}'_{\rho;q;C}(D_1, D_2)$.

10 Cohomology and duality

Clearly, our next objective, now that we have well-defined quasi-schemes associated to the difference operator construction, should be to show that the quasi-scheme behaves like a smooth surface. More precisely, we would like to show that $X_{\rho;q;C}$ is a "noncommutative smooth proper surface" in the sense of [14, Defn. 3.2]. The bulk of these axioms relate to the behavior of the Ext functor in $\operatorname{coh} X_{\rho;q;C}$, which will thus occupy our attention for the moment. As a special case, we define $\Gamma(M) := \operatorname{Hom}(\mathcal{O}_X, M)$, and $H^i(M) := \operatorname{Ext}^i(\mathcal{O}_X, M)$. (We should really denote this by Ext^i_X, but will suppress the surface from the notation, and only note when the Ext functor is being computed in some other (derived) category, typically $D^b \operatorname{coh} C$)

We first need to consider a few functors between our categories. We have already discussed the functor $\bar{\theta}$ with its associated natural transformation $T : \bar{\theta} \to \mathrm{id}$, as well as the functors $_(D)$ corresponding to twisting by a line bundle. We will also need a family of functors relating to the blowup structure. For convenience, let X_l denote the l-th noncommutative blowup in our sequence; that is,

$$X_l := X_{\eta,x_0,\dots,x_l;q;C}. \tag{10.1}$$

We include $X_0 = X_{\eta,x_0;q;C}$ as well as X_{-1}, the noncommutative $\mathbb{P}^2$. (We could also consider the analogue of F_0/F_2.)

Then Van den Bergh [64] gives us a pair of adjoint functors for each $0 \le l \le m$:

$$\begin{aligned} \alpha_{l*} &: \operatorname{qcoh} X_l \to \operatorname{qcoh} X_{l-1} \\ \alpha_l^* &: \operatorname{qcoh} X_{l-1} \to \operatorname{qcoh} X_l \end{aligned} \tag{10.2}$$

which preserve noetherian objects and structure sheaves. (And the expected exactness holds: α_{l*} is left-exact, and α_l^* is right-exact.) The first is quite straightforward in our terms, as it is simply the restriction (after saturation) to the appropriate $\mathbb{Z}^{l+1}$-algebra; the second is then the natural induction functor. We should note, however, that each functor comes in a natural 1-parameter family, for the simple reason that we could have restricted to a different coset of $\mathbb{Z}^{l+1}$. It is notationally convenient that conjugation by $\bar{\theta}$ moves between these alternate functors, and thus we need not introduce a special notation. The one exception is that we define (for reasons which will become apparent)

$$\alpha_l^! = \bar{\theta}^{-1} \alpha_l^* \bar{\theta}, \tag{10.3}$$

bearing in mind that the two copies of $\bar{\theta}$ are acting on different categories, so are not really the same functor. Note that $\alpha_l^!$ is right-exact (inherited from α_l^*); in addition, since α_l^* has a left-derived functor, so does $\alpha_l^!$. (The left-derived functor exists since line bundles are flat for α_l^* and $\alpha_l^!$.) In fact, Van den Bergh shows that α_l^* (and thus $\alpha_l^!$) has homological dimension 1, and similarly for α_{l*}.

Another key point is that Van den Bergh shows that $\bar{\theta}^{-1}$ is relatively ample for α_{l*}; that is, for any coherent sheaf M, there is a bound K_l such that for $k_l \ge K_l$, $\bar{\theta}^{-k_l} M$ is acyclic for α_{l*}, and $\alpha_l^* \alpha_{l*} \bar{\theta}^{-k_l} M \to M$ is surjective. This inductively gives a sequence $k_l, \dots, k_1$ such that

$$R\alpha_{1*} \bar{\theta}^{-k_1} R\alpha_{2*} \bar{\theta}^{-k_2} \cdots R\alpha_{l*} \bar{\theta}^{-k_l} M \tag{10.4}$$

is a sheaf on X_0, and thus (since ruled surfaces satisfy the Chan-Nyman axioms, [14]) there is a line bundle D on X_0 such that

$$R\operatorname{Hom}(\mathcal{O}_X(-D), R\alpha_{1*} \bar{\theta}^{-k_1} R\alpha_{2*} \bar{\theta}^{-k_2} \cdots R\alpha_{l*} \bar{\theta}^{-k_l} M) \tag{10.5}$$

is supported in cohomological degree 0. Since the functor being applied to M is a composition of functors with left adjoints, we may instead apply the adjoints to the line bundle, and thus find that there exists D on X_l such that $R\operatorname{Hom}(\mathcal{O}_X(-D), M)$ is supported in degree 0. (Moreover, such D exists in any translate of the universal nef cone.)

Define a sheaf $\mathcal{O}_{e_l}(-1) \cong \mathcal{O}_X(e_l)/\mathcal{O}_X$.

Lemma 10.1. *For any divisor D such that $X_{l-1}(D) = D$, there is a short exact sequence*

$$0 \to \mathcal{O}_X(D) \to \mathcal{O}_X(D + e_l) \to \mathcal{O}_{e_l}(-1) \to 0. \tag{10.6}$$

Proof. Up to scalars, there is a unique nonzero morphism $\mathcal{O}_X(D) \to \mathcal{O}_X(D+e_l)$, and that morphism induces an isomorphism

$$\operatorname{Hom}(\mathcal{O}_X(-D'), \mathcal{O}_X(D)) \to \operatorname{Hom}(\mathcal{O}_X(-D'), \mathcal{O}_X(D + e_l)) \tag{10.7}$$

for all $D' = X_{l-1}(D')$. It follows that if M is the cokernel, then $R\operatorname{Hom}(\mathcal{O}_X(-D'), M) = 0$ for all D', and thus $R\alpha_{l*}M = 0$. By [64, Thm. 8.4.1], this implies that M is isomorphic to a power of $\mathcal{O}_{e_l}(-1)$. Since $\dim \operatorname{Hom}(\mathcal{O}_X(-D' + e_l), M) = 1$ for sufficiently ample $D' = X_{l-1}(D')$, that power must be 1 as required. □

Proposition 10.2. *On coherent sheaves, the derived functor $L\alpha_l^!$ is right adjoint to the derived functor $R\alpha_{l*}$.*

Proof. We first claim that there is a natural isomorphism

$$R\alpha_{l*}L\alpha_l^! N \cong N \tag{10.8}$$

for any coherent sheaf N. This immediately reduces to the case that $N \cong \mathcal{O}_X(D)$, and thus to showing that

$$R\alpha_{l*}\mathcal{O}_X(D + e_l) \cong \mathcal{O}_X(D). \tag{10.9}$$

For any D', we have

$$\begin{aligned} R\operatorname{Hom}(\mathcal{O}_X(-D'), R\alpha_{l*}\mathcal{O}_X(D + e_l)) &\cong R\operatorname{Hom}(L\alpha_l^*\mathcal{O}_X(-D'), \mathcal{O}_X(D + e_l)) \\ &\cong R\operatorname{Hom}(\mathcal{O}_X(-D'), \mathcal{O}_X(D + e_l)), \end{aligned} \tag{10.10}$$

and thus the claim follows from the fact that for sufficiently ample D', $\mathcal{S}'_{\rho;q;C}(-D - e_l, D')$ is saturated and equal to $\mathcal{S}'_{\rho;q;C}(-D, D')$.

Now, to prove the Proposition, we need to show that for any coherent sheaves M, N, there is a natural isomorphism

$$R\operatorname{Hom}(M, L\alpha_l^! N) \cong R\operatorname{Hom}(R\alpha_{l*}M, N). \tag{10.11}$$

By [64, Thm. 8.4.1], there is a natural distinguished triangle

$$E \to L\alpha_l^* R\alpha_{l*}M \to M \to \tag{10.12}$$

where E is quasi-isomorphic to a (bounded) complex with entries of the form $\mathcal{O}_{e_l}(-1)^n$ and all morphisms 0. Now,

$$R\operatorname{Hom}(L\alpha_l^* R\alpha_{l*}M, L\alpha_l^! N) \cong R\operatorname{Hom}(R\alpha_{l*}M, R\alpha_{l*}L\alpha_l^! N) \cong R\operatorname{Hom}(R\alpha_{l*}M, N), \tag{10.13}$$

and thus it remains only to show that

$$R\operatorname{Hom}(\mathcal{O}_{e_l}(-1), L\alpha_l^! N) = 0. \tag{10.14}$$

The presentation $0 \to \mathcal{O}_X \to \mathcal{O}_X(e_l) \to O_{e_l}(-1) \to 0$ allows us to rephrase this as an isomorphism

$$R\operatorname{Hom}(\mathcal{O}_X, L\alpha_l^! N) \cong R\operatorname{Hom}(\mathcal{O}_X(e_l), L\alpha_l^! N) \tag{10.15}$$

or equivalently

$$R\operatorname{Hom}(L\alpha^*\mathcal{O}_X, L\alpha_l^! N) \cong R\operatorname{Hom}(\bar{\theta} L\alpha_l^* \bar{\theta}^{-1}\mathcal{O}_X, L\alpha_l^! N) \tag{10.16}$$

But this becomes

$$R\operatorname{Hom}(\mathcal{O}_X, R\alpha_* L\alpha_l^! N) \cong R\operatorname{Hom}(\bar{\theta}^{-1}\mathcal{O}_X, R\alpha_* L\alpha_l^* \bar{\theta}^{-1} N) \tag{10.17}$$

and thus reduces to the isomorphism $R\alpha_* L\alpha_l^! N \cong N$ already established. □

Remark. In contrast to the adjunction between α^* and α_*, this adjunction only holds at the level of derived categories; α_* is left exact, so cannot have a right adjoint as a functor of abelian categories.

Corollary 10.3. *There is a natural distinguished triangle*

$$M \to L\alpha_l^! R\alpha_{l*} M \to R\operatorname{Hom}(\mathcal{O}_{e_l}(-1), M) \otimes_k^{\mathbf{L}} \mathcal{O}_{e_l}(-1)[1] \to . \tag{10.18}$$

Proof. Applying $R\alpha_*$ to the natural morphism $M \to L\alpha^! R\alpha_{l*} M$ gives the identity on $R\alpha_{l*}M$, and thus the cone is in the image of $\otimes_k^{\mathbf{L}} \mathcal{O}_{e_l}(-1)$. Applying $R\operatorname{Hom}(\mathcal{O}_{e_l}(-1), _)$ then tells us which complex of vector spaces to take. □

This gives us the following weak form of Serre duality.

Proposition 10.4. *For any coherent sheaf M, and any divisor class D, there is a natural isomorphism*

$$R\operatorname{Hom}(M, \bar{\theta}\mathcal{O}_X(D)[2]) \cong R\operatorname{Hom}(\mathcal{O}_X(D), M)^*, \tag{10.19}$$

functorially in M.

Proof. By twisting X as necessary, we may assume $D = 0$. We also observe that this is known [37, 14] for X_0. We may thus assume by induction that it is known on the blown down surface, and proceed as follows:

$$\begin{aligned} R\operatorname{Hom}(M, \bar{\theta}\mathcal{O}_X[2]) &\cong R\operatorname{Hom}(M, \alpha_m^! \bar{\theta}\mathcal{O}_X[2]) \\ &\cong R\operatorname{Hom}(R\alpha_{m*}M, \bar{\theta}\mathcal{O}_X[2]) \\ &\cong R\operatorname{Hom}(\mathcal{O}_X, R\alpha_{m*}M)^* \\ &\cong R\operatorname{Hom}(\mathcal{O}_X, M)^*. \end{aligned} \tag{10.20}$$

□

Corollary 10.5. *For any divisor class D, if $D \cdot f > -2$, then $\operatorname{Ext}^2(\mathcal{O}_X, \mathcal{O}_X(D)) = 0$.*

Proof. This follows immediately from duality and the fact that $\mathcal{S}'_{\rho;q;C}(0, D) = 0$ if $D \cdot f < 0$. □

Corollary 10.6. *For any coherent sheaves M, N, $\operatorname{Ext}^p(M, N)$ is finite-dimensional for $p \geq 0$, and 0 for $p > 2$.*

Proof. We first show that there is a uniform bound on the Ext degree for any pair of coherent sheaves. We use the blowup structure and proceed by induction (the base case being known by [14]). Applying $R\operatorname{Hom}(M, _)$ to the distinguished triangle of Corollary 10.3 gives a distinguished triangle

$$R\operatorname{Hom}(M, \mathcal{O}_{e_m}(-1)) \otimes_k^{\mathbf{L}} R\operatorname{Hom}(\mathcal{O}_{e_m}(-1), N) \to R\operatorname{Hom}(M, N) \to R\operatorname{Hom}(R\alpha_* M, R\alpha_* N) \to \tag{10.21}$$

Now, $R\alpha_* M$ and $R\alpha_* N$ are bounded coherent complexes, and thus that term of the distinguished triangle is a uniformly bounded coherent complex by induction. The distinguished triangle of Corollary 10.3, as well as the analogue for $L\alpha^*$, tell us that $R\operatorname{Hom}(\mathcal{O}_{e_m}(-1), N)$ and $R\operatorname{Hom}(M, \mathcal{O}_{e_m}(-1))$ are uniformly bounded and coherent, respectively, and thus so is their tensor product. It thus follows that $R\operatorname{Hom}(M, N)$ is a uniformly bounded coherent complex.

It remains only to show that $\operatorname{Ext}^p(M, N) = 0$ for $p > 2$. Let q be the minimum bound such that $\operatorname{Ext}^p(M, N) = 0$ for $p > q$ and all coherent M, N, and consider a short exact sequence

$$0 \to N' \to \mathcal{O}_X(-D)^n \to N \to 0 \tag{10.22}$$

Weak Serre duality tells us that $\operatorname{Ext}^p(M, \mathcal{O}_X(-D)) \cong \operatorname{Ext}^{2-p}(\bar{\theta}^{-1}\mathcal{O}_X(-D), M)^*$, and thus for $p > 2$, $\operatorname{Ext}^p(M, \mathcal{O}_X(-D)) = 0$. It thus follows that $\operatorname{Ext}^p(M, N) = 0$ for $p > \max(q-1, 2)$, so that $q \leq \max(q-1, 2)$, i.e., $q \leq 2$. □

Proposition 10.7. *Any complex M of quasicoherent sheaves with bounded coherent cohomology is quasi-isomorphic to a bounded complex in which each term is a sum of line bundles; moreover, those line bundles may be restricted to a finite set that works for all M.*

Proof. This is known for X_0; indeed, we may take the line bundles to be $\mathcal{O}_X$, $\mathcal{O}_X(-f)$, $\mathcal{O}_X(-s-f)$, $\mathcal{O}_X(-s-2f)$. Thus as usual we reduce to showing that this is preserved under blowing up. We use the distinguished triangle

$$L\alpha_m^* R\alpha_{m*} M \to M \to R\operatorname{Hom}_k(R\operatorname{Hom}(M, \mathcal{O}_{e_m}(-1)), \mathcal{O}_{e_m}(-1)) \to \tag{10.23}$$

Since $R\alpha_{m*} M$ has bounded coherent cohomology, it is quasi-isomorphic to a bounded complex of sums of line bundles, and this is preserved under $L\alpha_m^*$. Since the set of line bundles used can be fixed independently of M, we may in particular choose a divisor class $D_m = X_{m-1}(D_m)$ such that for any of those line bundles,

$$R\operatorname{Hom}(\mathcal{O}_X(-D_m), \mathcal{O}_X(D)) \quad \text{and} \quad R\operatorname{Hom}(\mathcal{O}_X(-D_m + e_m), \mathcal{O}_X(D)) \tag{10.24}$$

are supported in degree 0. We may express $\mathcal{O}_{e_m}(-1) = \mathcal{O}_X(-D_m + e_m)/\mathcal{O}_X(-D_m)$, and thus find that

$$R\operatorname{Hom}_k(R\operatorname{Hom}(M, \mathcal{O}_{e_m}(-1)), \mathcal{O}_{e_m}(-1)) \tag{10.25}$$

is quasi-isomorphic to a complex in which each term is a sum of $\mathcal{O}_X(-D_m)$ and $\mathcal{O}_X(-D_m + e_m)$. Since D_m was sufficiently ample, the map

$$R\operatorname{Hom}_k(R\operatorname{Hom}(M, \mathcal{O}_{e_m}(-1)), \mathcal{O}_{e_m}(-1)) \to L\alpha_m^* R\alpha_{m*} M[1] \tag{10.26}$$

corresponds to an actual map of complexes, and M is represented by the cone. □

Remark. Note that the argument works equally well when M and N are flat families over some base (including if the parameters $\rho; q; C$ themselves vary over the base).

This allows us to define an Euler characteristic

$$\chi(M,N) := \dim \mathrm{Hom}(M,N) - \dim \mathrm{Ext}^1(M,N) + \dim \mathrm{Ext}^2(M,N), \tag{10.27}$$

which is constant as M and N vary over flat families.

Proposition 10.8. *For any coherent sheaf M, there are integers $\mathrm{rank}(M)$, $\chi(M)$ and a divisor class $c_1(M)$ such that for any divisor class D,*

$$\chi(\mathcal{O}_X(-D), M) = \mathrm{rank}(M)\frac{D\cdot(D+C_m))}{2} + c_1(M)\cdot D + \chi(M). \tag{10.28}$$

Proof. Since the function $D \mapsto \chi(O_X(-D), M)$ is additive in short exact sequences, we see that for any exact sequence

$$0 \to M_1 \to M_2 \to M_3 \to 0, \tag{10.29}$$

if the Proposition holds for two of the terms, then it holds for the third, with

$$\begin{aligned} \mathrm{rank}(M_2) &= \mathrm{rank}(M_1) + \mathrm{rank}(M_3), \\ c_1(M_2) &= c_1(M_1) + c_1(M_3), \\ \chi(M_2) &= \chi(M_1) + \chi(M_3). \end{aligned} \tag{10.30}$$

Since any coherent sheaf M is represented by a bounded complex of line bundles, it thus suffices to prove the Proposition when M is itself a line bundle. Since every line bundle is a fiber of a flat family over the entirety of parameter space, we find that $\chi(\mathcal{O}_X(-D), \mathcal{O}_X(D'))$ is constant over parameter space. When $q = 1$, we find that

$$\chi(\mathcal{O}_X(-D), \mathcal{O}_X(D')) = \chi(\mathcal{O}_X(D+D')) = 1 + \frac{(D+D')\cdot(D+D'+C_m)}{2}, \tag{10.31}$$

and thus $\mathcal{O}_X(D')$ satisfies the given claim with

$$\begin{aligned} \mathrm{rank}(\mathcal{O}_X(D')) &= 1 \\ c_1(\mathcal{O}_X(D')) &= D' \\ \chi(\mathcal{O}_X(D')) &= 1 + \frac{D'\cdot(D'+C_m)}{2}. \end{aligned} \tag{10.32}$$

□

Of course, the notation is consistent with the usual notation in the commutative case; we will thus refer to $\mathrm{rank}(M)$, $c_1(M)$ and $\chi(M)$ as the "rank", "first Chern class", and "Euler characteristic" of M even in the noncommutative case. We will also refer to the tuple $(\mathrm{rank}(M), c_1(M), \chi(M))$ as the "numerical invariants" of M. Also, since they are defined using $\chi(M,N)$, these quantities themselves must be constant in flat families.

Corollary 10.9. *The product map $(\mathrm{rank}, c_1, \chi) : K_0(\mathrm{coh}\, X_m) \to \mathbb{Z}^{m+4}$ is an isomorphism.*

Proof. The usual induction tells us that $K_0(\mathrm{coh}\, X_m)$ is generated by the classes

$$[\mathcal{O}_X], [\mathcal{O}_X(-f)], [\mathcal{O}_X(-s-f)], [\mathcal{O}_X(-s-2f)], [\mathcal{O}_X(e_1)], \ldots, [\mathcal{O}_X(e_m)]. \tag{10.33}$$

It is straightforward to see that their images in $\mathbb{Z}^{m+4}$ form a basis, and thus they are linearly independent classes in $K_0(\mathrm{coh}\, X_m)$, giving the desired isomorphism. □

Corollary 10.10. *For any coherent sheaves M, N,*

$$\chi(M,N) = -\operatorname{rank}(M)\operatorname{rank}(N) + \operatorname{rank}(M)\chi(N) + \chi(M)\operatorname{rank}(N) - c_1(M)\cdot(c_1(N) + \operatorname{rank}(N)C_m)$$

Proof. The function $\chi(M,N)$ is a bilinear form on the Grothendieck group; since the latter has a basis of line bundles, we need simply verify this when both M and N are line bundles. □

Comparing the two sides of weak Serre duality gives the following.

Proposition 10.11. *For any coherent sheaf M,*

$$\operatorname{rank}(\bar{\theta}M) = \operatorname{rank}(M) \tag{10.34}$$
$$c_1(\bar{\theta}M) = c_1(M) - \operatorname{rank}(M)C_m \tag{10.35}$$
$$\chi(\bar{\theta}M) = \chi(M) - c_1(M)\cdot C_m. \tag{10.36}$$

Remark. And, of course, $\chi(M,N) = \chi(N,\bar{\theta}M)$ in general, since it holds for line bundles.

It is also straightforward to see how these numerical invariants are affected by the twisting and adjoint operations. Let $M(D)$ be the image of the twist functor taking $\mathcal{O}_X(D') \to \mathcal{O}_X(D+D')$ (and acting in the appropriate way on ρ).

Proposition 10.12. *For any coherent sheaf M,*

$$\begin{aligned}\operatorname{rank}(M(D)) &= \operatorname{rank}(M)\\ c_1(M(D)) &= c_1(M) + \operatorname{rank}(M)D\\ \chi(M(D)) &= \chi(M) + c_1(M)\cdot D + \operatorname{rank}(M)\frac{D\cdot(D+C_m)}{2},\end{aligned}$$

Since $\mathcal{O}_X(D)|_C \cong \mathcal{L}(\mathfrak{D}_{\rho;q;C}(D))$, we may also easily compute how $|_C^{\mathbf{L}}$ acts on Grothendieck groups.

Proposition 10.13. *Let M be a coherent sheaf on $X_{\rho;q;C}$, with $M|_C^{\mathbf{L}}$ the corresponding complex in* $\operatorname{coh} C$*. Then*

$$\operatorname{rank}(M|_C^{\mathbf{L}}) = \operatorname{rank}(M) \tag{10.37}$$
$$\deg(M|_C^{\mathbf{L}}) = c_1(M)\cdot C_m \tag{10.38}$$
$$c_1(M|_C^{\mathbf{L}}) = q^{\chi(M)-\operatorname{rank}(M)}\rho(c_1(M)). \tag{10.39}$$

Note, in particular, that we can recover ρ and q from the restriction map $K_0(\operatorname{coh} X) \to K_0(\operatorname{coh} C)$.

Corollary 10.14. *If M is disjoint from C, then M is a rank 0 sheaf transverse to C, and T_M is an isomorphism.*

Proof. If $M|_C = 0$, then $\operatorname{rank}(M) = \operatorname{rank}(M|_C^{\mathbf{L}}) \le 0$, and thus M must be have rank 0. In particular, $\ker(T_M)$ is 0-dimensional, and if it is nonzero, then $c_1(M)\cdot C_m < 0$. This implies that $\chi(\mathcal{L}, M) > 0$ for all line bundles $\mathcal{L}$ on C and thus that there exists a morphism from M to some such line bundle, contradicting disjointness. Thus M is transverse to C, and since both the kernel and cokernel of T_M vanish, T_M is an isomorphism. □

Proposition 10.15. *If D is nef with $D\cdot C_m > 0$, then $\operatorname{Ext}^i(\mathcal{O}_X, \mathcal{O}_X(D)) = 0$ for $i > 0$.*

Proof. For $i \geq 2$, this follows from weak Serre duality, at which point it follows for $i = 1$ from flatness:

$$\dim \operatorname{Hom}(\mathcal{O}_X, \mathcal{O}_X(D)) = \chi(\mathcal{O}_X, \mathcal{O}_X(D)), \tag{10.40}$$

and thus $\dim \operatorname{Ext}^1(\mathcal{O}_X, \mathcal{O}_X(D)) = 0$. □

With this, it is relatively straightforward to obtain semicontinuity.

Lemma 10.16. *Let X/S be a family of quasi-schemes of the form $X_{\rho;q;C}$ (with fixed m) over the locally noetherian scheme S, and let $M \in D^b \operatorname{coh} X$. Then for all $i \in \mathbb{Z}$, $s \mapsto \dim \operatorname{Ext}^i(\mathcal{O}_X, M|_s^{\mathbf{L}})$ is an upper semicontinuous function of s.*

Proof. There is a fixed finite collection of divisor classes D_i such that for any divisor class D, we may represent M as a bounded complex in which each term is a sum

$$\bigoplus_i \mathcal{O}_X(D + D_i) \boxtimes_S F_i \tag{10.41}$$

with F_i a sheaf on S. Indeed, this is just the relative version of Proposition 10.7, applied to the twisted complex $M(-D)$. If $D + D_i$ is ample for every D, this gives a bounded complex of sheaves on S representing $R\operatorname{Hom}(\mathcal{O}_X, M|_s^{\mathbf{L}})$, and we can then apply the usual argument to a locally free resolution of this complex. □

Remark. We can also apply this to $\operatorname{Ext}^i(M, N)$ for M, $N \in D^b \operatorname{coh} X$, as long as M has a representation as a bounded coherent complex in which each term is a sum of tensor products of line bundles on X with *locally free* sheaves on S. Over an integral (or locally integral) base, this can always be done when $M \in \operatorname{coh} X$; the coefficient sheaves arising in the construction can themselves be computed as Ext sheaves to which semicontinuity applies, and if we twist by a sufficiently ample bundle, we can arrange for there to be only one nonvanishing Ext group in each case. At that point, Euler characteristic considerations tell us that the dimension is constant, and the analogue of Grauert's theorem tells us that the sheaves are locally free.

One use of the above linear functionals on the Grothendieck group is to define a "size" for coherent sheaves.

Definition 10.1. A nonzero coherent sheaf M is 2-*dimensional* if $\operatorname{rank}(M) \neq 0$, 1-*dimensional* if $\operatorname{rank}(M) = 0$ and $c_1(M) \neq 0$, and 0-*dimensional* if $\operatorname{rank}(M) = 0$ and $c_1(M) = 0$. If no subsheaf of M has smaller dimension, we say that M is a *pure* d-dimensional sheaf.

Proposition 10.17. *A 2-dimensional coherent sheaf has positive rank; a 0-dimensional coherent sheaf has positive Euler characteristic. Finally, if M is 1-dimensional, then for any ample divisor class D, $D \cdot c_1(M) > 0$.*

Proof. For any ample divisor D, we must have $\operatorname{Hom}(\mathcal{O}_X(-rD), M) > 0$ for $r \gg 0$. This implies $\operatorname{rank}(M) > 0$ or $\chi(M) > 0$ in the 2- or 0-dimensional cases respectively, while in the 1-dimensional case implies $D \cdot c_1(M) > 0$. □

Corollary 10.18. *For any coherent d-dimensional sheaf M, there is a bound on the length of a descending chain $M = M_0 \supset M_1 \supset \cdots$ of subsheaves such that all quotients are d-dimensional.*

Proof. Fix an ample divisor D, and consider the following identity of polynomials:

$$\chi(\mathcal{O}_X(-rD), M) = \chi(\mathcal{O}_X(-rD), M_0/M_1) + \chi(\mathcal{O}_X(-rD), M_1/M_2) + \cdots + \chi(\mathcal{O}_X(-rD), M_l) \quad (10.42)$$

for all l. By assumption, all of the terms on either side of this expression are polynomials of the same degree with positive leading coefficient. Since the leading coefficients are at worst half-integers, the number of terms on the right can be bounded in terms of the leading coefficient on the left; at most $\operatorname{rank}(M)$, $c_1(M) \cdot D$ or $\chi(M)$, depending on d. □

This leads to the following analogue of the usual definition of nef in the commutative setting.

Corollary 10.19. *A divisor class D is nef iff for all 1-dimensional sheaves M, $D \cdot c_1(M) \geq 0$.*

Proof. If M is 1-dimensional, then $D \cdot c_1(M) > 0$ for all ample divisors; since any nef divisor is a limit of rational multiples of ample divisors, we must at least have a weak inequality when D is nef.

Now, let D be any divisor class having nonnegative intersection with the Chern class of every 1-dimensional sheaf. The sheaves $\mathcal{O}_C$ and $\mathcal{O}_{e_m}(-1)$ have Chern classes C_m and e_m respectively, and thus $D \cdot C_m, D \cdot e_m \geq 0$; we similarly have $D \cdot f \geq 0$ using a sheaf of the form $\mathcal{O}_X/\mathcal{O}_X(-f)$. If $D \cdot \alpha \geq 0$ for all simple roots α, then D is universally nef, and we are done. Otherwise, choose a simple root α with $D \cdot \alpha < 0$. If $\rho(\alpha) \in q^{\mathbb{Z}}$, then there is a divisor class D' such that

$$\hat{\mathcal{S}}'_{\rho;q;C}(D', D' + \alpha) \neq 0, \quad (10.43)$$

and the cokernel of such a morphism will be a 1-dimensional sheaf with Chern class α, giving a contradiction. Thus $\rho(\alpha) \notin q^{\mathbb{Z}}$, so that we may perform the corresponding simple reflection. Repeating this will either force $D \cdot f < 0$ (a contradiction) or put D into the universal nef cone, meaning that the original divisor class was nef. □

Remark. Note, in particular, that this implies that the nef divisors are the integer points of a cone, as our terminology had suggested.

With this in mind, we call a divisor class on X_m *effective* if it is the Chern class of a 1-dimensional sheaf, or equivalently if it is in the dual to the nef cone. The set of effective divisor classes is clearly the monoid generated by the Chern classes of *irreducible* sheaves (1-dimensional sheaves such that any subsheaf has the same Chern class). Moreover, we can identify some such irreducible Chern classes: -1-curves, i.e., those with $e^2 = -1$, $e \cdot C_m = 1$, and -2-curves, those with $e^2 = -2$, $e \cdot C_m = 0$. The prototypical example of a -1-curve is of course e_m, and it follows as in the commutative case that any -1-curve is related by an admissible element of $W(E_{m+1})$ to this case; in particular, any -1-curve can be blown down. Similarly, any -2-curve is admissibly equivalent to a simple root.

We can further argue as in [47, Cor. 4.5] to obtain the following.

Corollary 10.20. *If $m > 0$, then the effective cone is generated by -1-curves, -2-curves, and C_m.*

We also have the following.

Lemma 10.21. *Given any divisor class $D \in \operatorname{Pic}(X_m)$, there are only finitely many divisor classes D' such that $D - D'$ and D' are both effective.*

Proof. Although the effective cone need not be finitely generated, it is contained in the cone generated by simple roots, e_m, and C_m (since the nef cone contains the universal nef cone). If D_a is a universally ample divisor, then D_a has positive intersection with the generators of this larger cone, and thus there are only finitely many D' in the larger cone such that $D - D'$ is also in the larger cone. □

Corollary 10.22. *A divisor class D is ample iff for all 1-dimensional sheaves M, $D \cdot c_1(M) > 0$.*

Proof. Indeed, this is precisely the same as saying that D is in the interior of the nef cone. □

Corollary 10.23. *A subquotient of a d-dimensional sheaf has dimension at most d.*

Proof. If $\operatorname{rank}(M) = 0$, then for any subsheaf M', we have $\operatorname{rank}(M') + \operatorname{rank}(M/M') = 0$, and thus both must vanish. Similarly, if $\operatorname{rank}(M) = 0$ and $c_1(M) = 0$, then $c_1(M') = c_1(M/M') = 0$. □

Proposition 10.24. *Any nonzero subsheaf of $\mathcal{O}_X$ has rank 1.*

Proof. Let M be a nonzero subsheaf of $\mathcal{O}_X$. Then there exists a divisor class D such that $\operatorname{Hom}(\mathcal{O}_X(-D), M) \neq 0$, and composing gives a morphism $\mathcal{O}_X(-D) \to \mathcal{O}_X$. Since $\hat{\mathcal{S}}'_{\rho;q;C}$ is a domain, any such morphism is injective, and thus $M \supset \mathcal{O}_X(-D)$. It follows that

$$1 = \operatorname{rank}(\mathcal{O}_X(-D)) \leq \operatorname{rank}(M) \leq \operatorname{rank}(\mathcal{O}_X) \leq 1. \tag{10.44}$$

□

Remark. It follows easily that any sum of line bundles is a pure 2-dimensional sheaf.

Corollary 10.25. *If M is 1-dimensional, then for any divisor class D, $\operatorname{Hom}(M, \mathcal{O}_X(D)) = 0$ and $\operatorname{Ext}^2(\mathcal{O}_X(-D), M) = 0$.*

Proof. The two claims are equivalent by weak Serre duality, so it suffices to show that for all D, $\operatorname{Hom}(M, \mathcal{O}_X(D)) = 0$. By twisting, we may reduce to the case $D = 0$. The image of any nonzero morphism $M \to \mathcal{O}_X$ must be 1-dimensional, as a quotient of M, and 2-dimensional, as a subsheaf of $\mathcal{O}_X$, and this gives a contradiction. □

Proposition 10.26. *If q is nontorsion, then for any 0-dimensional sheaf M, the morphism $T_M^k : \bar{\theta}^k M \to M$ is 0 for some $k > 0$.*

Proof. Since $\operatorname{rank}(M) = 0$ and $c_1(M) = 0$, we compute that $\operatorname{rank}(M|_C^{\mathbf{L}}) = 0$ and $c_1(M|_C^{\mathbf{L}}) = q^{\chi(M)}$. If $T : \theta M \to M$ were surjective, then the kernel would need to be a nontrivial (since $q^{\chi(M)} \neq 1$) sheaf of rank 0 and degree 0 on C, a contradiction. We thus have an exact sequence

$$0 \to M' \to M \to M|_C \to 0, \tag{10.45}$$

where, by induction in $\chi(M)$, $\bar{\theta}^k M' \to M'$ is 0 for some $k \geq 0$, forcing $\bar{\theta}^{k+1} M \to M$ to be 0. □

Proposition 10.27. *If M is 0-dimensional, then $\operatorname{Ext}^p(\mathcal{O}_X, M) = 0$ for $p > 0$.*

Proof. Since the claim is preserved under extensions, we may restrict our attention to the case that M has no proper subsheaves. In particular, either $T_M = 0$ or $\bar{\theta} M \cong M$, with the latter only occurring in the case that q is torsion. In the former case, M is the structure sheaf of a point of C, and may thus be expressed as a quotient of some acyclic line bundle, which implies the desired result by dimension considerations. On the other hand, if $\bar{\theta} M \cong M$, then we immediately find that $R\alpha_{m*} M$ is a sheaf, which must itself be 0-dimensional, letting us reduce to the Hirzebruch surface case, where it is known [14]. □

Although we have seen that weak Serre duality is already quite powerful, we would of course like to have the full form of duality. The main difficulty we encounter is that the above construction of the pairing corresponding to weak Serre duality is not functorial in the line bundle, which prevents us from using the fact that we can represent every sheaf as a bounded complex of sums of line bundles.

A related difficulty is that, although weak Serre duality gives us maps

$$\mathrm{Ext}^2(\mathcal{O}_X(D), \bar{\theta}\mathcal{O}_X(D)) \cong k, \tag{10.46}$$

these maps are not truly canonical. Now, the map $T : \bar{\theta}\mathcal{O}_X(D) \to \mathcal{O}_X(D)$ is injective, and thus we have a morphism

$$\mathrm{Ext}^1(\mathcal{O}_X(D), \mathcal{O}_X(D)|_C) \to \mathrm{Ext}^2(\mathcal{O}_X(D), \bar{\theta}\mathcal{O}_X(D)). \tag{10.47}$$

Since $\mathrm{Ext}^i(\mathcal{O}_X(D), \mathcal{O}_X(D)) = 0$ for $i > 0$, this map is in fact an isomorphism. We also have a *canonical* isomorphism

$$\mathrm{Ext}^1(\mathcal{O}_X(D), \mathcal{O}_X(D)|_C) \cong \mathrm{Ext}^1_C(\mathcal{O}_X(D)|_C, \mathcal{O}_X(D)|_C) \cong H^1_C(\mathcal{O}_C), \tag{10.48}$$

and thus we see that we can reduce the choice of isomorphism

$$\mathrm{Ext}^2(\mathcal{O}_X(D), \bar{\theta}\mathcal{O}_X(D)) \cong k \tag{10.49}$$

to a choice of nonzero holomorphic differential on C. With this in mind, we define a new functor θ by

$$\theta M := \bar{\theta}M \otimes_k H^0_C(\omega_C), \tag{10.50}$$

and observe that the above calculation gives *natural* isomorphisms

$$\mathrm{Ext}^2(\mathcal{O}_X(D), \theta\mathcal{O}_X(D)) \to k. \tag{10.51}$$

Lemma 10.28. *The derived category $D^b \operatorname{coh} X_{\rho;q;C}$ is generated by an exceptional collection.*

Proof. Indeed, the sequence

$$\mathcal{O}_X(e_m)/\mathcal{O}_X, \mathcal{O}_X(e_{m-1})/\mathcal{O}_X, \dots \mathcal{O}_X(e_1)/\mathcal{O}_X, \mathcal{O}_X(-s-2f), \mathcal{O}_X(-s-f), \mathcal{O}_X(-f), \mathcal{O}_X \tag{10.52}$$

is an exceptional collection, and we have already seen that it generates. □

Remark. Another possibility is

$$\mathcal{O}_X(e_m)/\mathcal{O}_X, \mathcal{O}_X(e_{m-1})/\mathcal{O}_X, \dots \mathcal{O}_X(e_1)/\mathcal{O}_X, \mathcal{O}_X(s)/\mathcal{O}_X, \mathcal{O}_X(-2s-2f), \mathcal{O}_X(-s-f), \mathcal{O}_X, \tag{10.53}$$

corresponding to blowing up $\mathbb{P}^2$ rather than F_1.

Theorem 10.29. *The functor $M \mapsto \theta M[2]$ is a Serre functor on $D^b \operatorname{coh} X_{\rho;q;C}$. That is, for any coherent sheaves M, N, there are functorial isomorphisms*

$$\mathrm{Ext}^i(M, N) \cong \mathrm{Ext}^{2-i}(N, \theta M)^*. \tag{10.54}$$

Proof. Since $D^b \operatorname{coh} X_{\rho;q;C}$ has bounded and finite Ext groups, and is generated by an exceptional collection, it follows from [10] that there *is* a Serre functor on $D^b \operatorname{coh} X_{\rho;q;C}$. Write this Serre functor as $\hat{\theta}[2]$ (where $\hat{\theta}$ is in principle only a *derived* equivalence). Weak Serre duality gives

$$\mathrm{Ext}^i(M, \theta\mathcal{O}_X(D)) \cong \mathrm{Ext}^{2-i}(\mathcal{O}_X(D), M)^* \cong \mathrm{Ext}^i(M, \hat{\theta}\mathcal{O}_X(D)) \tag{10.55}$$

for any sheaf M, functorially in M; taking $M = \mathcal{O}_X(-D')$ for sufficiently ample D' implies that $\hat{\theta}\mathcal{O}_X(D)$ is not only a sheaf, but isomorphic to $\theta\mathcal{O}_X(D)$. Since both morphisms are canonical in the case $M = \theta\mathcal{O}_X(D)$, we conclude that we have a natural isomorphism $\theta \cong \hat{\theta}$ as required. □

We also want to consider another useful form of duality, analogous to the functor $R\mathcal{H}om(_, \omega_X)$ in the commutative case. Since the formal adjoint induces an isomorphism $(\hat{\mathcal{S}}'_{\rho;q;C})^{\mathrm{op}} \cong \hat{\mathcal{S}}'_{\rho;q;C}$, it induces a contravariant equivalence ad between the subcategories of $\operatorname{coh} X_{\rho;q;C}$ and $\operatorname{coh} X_{\rho;1/q;C}$ induced by the line bundles. We fix a scaling for this functor by insisting that $\operatorname{ad} \mathcal{O}_X(D) \cong \theta \mathcal{O}_X(-D)$.

Proposition 10.30. *The functor* ad *extends to a contravariant equivalence* $R\operatorname{ad} : D^b \operatorname{coh} X^{\mathrm{op}}_{\rho;q;C} \to D^b \operatorname{coh} X_{\rho;1/q;C}$*, with* $R\operatorname{ad} R\operatorname{ad} \cong \mathrm{id}$ *and* $R\operatorname{ad} \mathcal{O}_X(D) \cong \theta \mathcal{O}_X(-D)$*. Moreover,* $R\operatorname{ad} \alpha_m^* \cong \alpha_m^! R\operatorname{ad}$*,* $R\operatorname{ad} \theta \cong \theta^{-1} R\operatorname{ad}$*, and* $R\operatorname{ad}$ *acts on Chern classes as*

$$\operatorname{rank}(R\operatorname{ad} M) = \operatorname{rank}(M) \tag{10.56}$$

$$c_1(R\operatorname{ad} M) = -c_1(M) - \operatorname{rank}(M) C_m \tag{10.57}$$

$$\chi(R\operatorname{ad} M) = \chi(M). \tag{10.58}$$

Proof. If an extension $R\operatorname{ad}$ exists with the above properties, it must in particular satisfy

$$R\operatorname{Hom}_{X_{\rho;1/q;C}}(\mathcal{O}_X(-D), R\operatorname{ad} M) \cong R\operatorname{Hom}_{X_{\rho;q;C}}(M, \theta \mathcal{O}_X(D)) \tag{10.59}$$

for all divisor classes D. Now, the right-hand side certainly gives a well-defined functor from $D^b \operatorname{coh} X_{\rho;q;C}$ to $D^b(\mathcal{S}'_{\rho;q;C})^{\mathrm{op}}$, and thus after applying the formal adjoint and quotienting by torsion modules, we obtain a functor $R\operatorname{ad} : D^b \operatorname{coh} X^{\mathrm{op}}_{\rho;q;C} \to D \operatorname{qcoh} X_{\rho;1/q;C}$, with $R^0 \operatorname{ad}$ agreeing with ad on line bundles since ad is an equivalence.

Since

$$R\operatorname{Hom}_{X_{\rho;q;C}}(\mathcal{O}_X(D), \theta \mathcal{O}_X(D')) \tag{10.60}$$

is supported in degree 0 for $D' - C_m - D$ nef, so is

$$R\operatorname{Hom}_{X_{\rho;1/q;C}}(\mathcal{O}_X(-D'), R\operatorname{ad} \mathcal{O}_X(D)). \tag{10.61}$$

But this implies that $R\operatorname{ad} \mathcal{O}_X(D)$ is a sheaf. In particular, applying $R\operatorname{ad}$ to a bounded complex of finite sums of line bundles gives another such complex, and thus $R\operatorname{ad}$ preserves bounded coherent complexes in general. The remaining claims follow similarly by reduction to the line bundle case. □

Remark. It will be convenient to refer to $X_{\rho;1/q;C}$ as $\operatorname{ad} X_{\rho;q;C}$, especially in cases when we wish to suppress parameters.

This satisfies the usual compatibility condition relative to the inclusion of C in X.

Proposition 10.31. *For any sheaf* M *supported on* C*, there is a natural isomorphism*

$$R\operatorname{ad} M \cong R\mathcal{H}om_C(M, \omega_C)[-1]. \tag{10.62}$$

Proof. We have

$$\begin{aligned}
R\operatorname{Hom}_C(\mathcal{O}_X(-D)|_C, R\operatorname{ad} M) &\cong R\operatorname{Hom}_{\operatorname{ad} X}(\mathcal{O}_X(-D), R\operatorname{ad} M) \\
&\cong R\operatorname{Hom}_X(M, \theta \mathcal{O}_X(D)) \\
&\cong R\operatorname{Hom}_X(\mathcal{O}_X(D), M[2]) \\
&\cong R\operatorname{Hom}_C(\mathcal{O}_X(D)|_C, M[2]) \\
&\cong R\operatorname{Hom}_C(M, \omega_C \otimes \mathcal{O}_X(D))[-1] \\
&\cong R\operatorname{Hom}_C(\mathcal{O}_X(-D)|_C, \mathcal{H}om_C(M, \omega_C))[-1].
\end{aligned}$$

□

Comparing to the axioms of [14], we see that there is only one thing left to prove, namely that Quot schemes are countable unions of projective schemes. Of course, in the commutative setting, they are projective as long as we specify the Hilbert polynomial of the quotient.

The usual treatment of this in the literature relies on the theory of Castelnuovo-Mumford regularity. It is unclear whether there is an analogue of this theory for general very ample divisor classes on X_m. The main difficulty is that a very ample divisor class in principle need not correspond to an embedding in a noncommutative $\mathbb{P}^n$. The proof in [35, §14] of the main properties of Castelnuovo-Mumford regularity involves an induction via hyperplane sections, which would seem to be another issue, but in fact allows us to abstract away the embedding. The key point is that the cohomology of a restriction of a sheaf to a hyperplane can be expressed in terms of the hyper-Ext from an appropriate linear complex.

For an ample divisor class D_a, call a coherent sheaf (l, D_a)-*regular* if $\operatorname{Ext}^2(\mathcal{O}_X((l+2)D_a), M) = \operatorname{Ext}^1(\mathcal{O}_X((l+1)D_a), M) = 0$, and D_a-*regular* if it is $(0, D_a)$-regular.

Lemma 10.32. *Suppose D_a is an ample divisor class such that for all l there is a sheaf F_l with a presentation of the form*

$$0 \to \mathcal{O}_X((l-3)D_a) \to \mathcal{O}_X((l-2)D_a)^{c_2} \to \mathcal{O}_X((l-1)D_a)^{c_1} \to \mathcal{O}_X(lD_a)^{c_0} \to F_l \to 0. \tag{10.63}$$

Then any $(0, D_a)$-regular sheaf is globally generated and (l, D_a)-regular for all $l \geq 0$.

Proof. Suppose M is a $(0, D_a)$-regular sheaf. Since $\mathcal{O}_X(D_a) \subset \mathcal{O}_X(2D_a)$, we have a surjection $\operatorname{Ext}^2(\mathcal{O}_X(2D_a), M) \to \operatorname{Ext}^2(\mathcal{O}_X(D_a), M)$, and thus the latter also vanishes. Now, the truncated complex

$$0 \to \mathcal{O}_X \to \mathcal{O}_X(D_a)^{c_2} \to \mathcal{O}_X(2D_a)^{c_1} \tag{10.64}$$

represents a sheaf, and thus the hyper-Ext from the complex to M is supported in degrees 0, 1, 2. The E_1 page of the corresponding spectral sequence has terms

$$\begin{array}{ccccc} 0 & \longrightarrow & 0 & \longrightarrow & \operatorname{Ext}^2(\mathcal{O}_X, M) \\ \operatorname{Ext}^1(\mathcal{O}_X(2D_a)^{c_1}, M) & \longrightarrow & 0 & \longrightarrow & \operatorname{Ext}^1(\mathcal{O}_X, M) \\ \operatorname{Hom}(\mathcal{O}_X(2D_a)^{c_1}, M) & \longrightarrow & \operatorname{Hom}(\mathcal{O}_X(D_a)^{c_2}, M) & \longrightarrow & \operatorname{Hom}(\mathcal{O}_X, M) \end{array} \tag{10.65}$$

and thus M is acyclic. Since then $M(D_a)$ is D_a-regular, it remains only to show that M is globally generated. Here we use the spectral sequence for $R\operatorname{Hom}(F_2, M)$, where we find that the sheaves that would map to E^{03} on all later pages already vanish on the E_1 page, and thus we have an exact sequence

$$\operatorname{Hom}(\mathcal{O}_X, M)^{c_2} \to \operatorname{Hom}(\mathcal{O}_X(-D_a), M) \to E_\infty^{03} \to 0. \tag{10.66}$$

But $\operatorname{Ext}^3(F_2, M) = 0$ and thus $E_\infty^{03} = 0$, so we conclude that $\operatorname{Hom}(\mathcal{O}_X, M)^{c_2} \to \operatorname{Hom}(\mathcal{O}_X(-D_a), M)$ is surjective. That map factors through the natural map

$$\operatorname{Hom}(\mathcal{O}_X(-D_a), \mathcal{O}_X) \otimes_k \operatorname{Hom}(\mathcal{O}_X, M) \to \operatorname{Hom}(\mathcal{O}_X(-D_a), M), \tag{10.67}$$

which thus must also be surjective. It follows by induction that the natural map

$$\operatorname{Hom}(\mathcal{O}_X(-lD_a), \mathcal{O}_X) \otimes_k \operatorname{Hom}(\mathcal{O}_X, M) \to \operatorname{Hom}(\mathcal{O}_X(-lD_a), M) \tag{10.68}$$

is surjective for all $l \geq 0$. Indeed, the composition

$$\begin{aligned} &\operatorname{Hom}(\mathcal{O}_X(-(l+1)D_a), \mathcal{O}_X(-lD_a)) \otimes_k \operatorname{Hom}(\mathcal{O}_X(-lD_a), \mathcal{O}_X) \otimes_k \operatorname{Hom}(\mathcal{O}_X, M) \\ &\quad \to \operatorname{Hom}(\mathcal{O}_X(-(l+1)D_a), \mathcal{O}_X(-lD_a)) \otimes_k \operatorname{Hom}(\mathcal{O}_X(-lD_a), M) \\ &\quad \to \operatorname{Hom}(\mathcal{O}_X(-(l+1)D_a), M) \end{aligned} \tag{10.69}$$

is surjective, and factors through

$$\mathrm{Hom}(\mathcal{O}_X(-(l+1)D_a),\mathcal{O}_X)\otimes_k \mathrm{Hom}(\mathcal{O}_X,M)\to \mathrm{Hom}(\mathcal{O}_X(-(l+1)D_a),M). \tag{10.70}$$

Since D_a is ample, this implies that M is globally generated as required. □

Corollary 10.33. *Any ample divisor of the form $s+d'f$ on X_0 or X_0' satisfies the hypothesis of Lemma 10.32.*

Proof. Let $m=D_a^2-2>0$, and let $y_1,\dots,y_m$ be a sufficiently general m-tuple of points on C. If we blow up those points, then $D_a-e_1-\cdots-e_m$ is $W(E_{m+1})$-equivalent to f. We thus find that there is an acyclic complex of the form

$$\mathcal{O}_X(e_1+\cdots+e_m-D_a+D)\to\mathcal{O}_X(D)^2\to\mathcal{O}_X(D+D_a-e_1-\cdots-e_m) \tag{10.71}$$

for any divisor D on X_m. (This follows from the corresponding fact for f.) Taking the Yoneda product with the extension

$$0\to\mathcal{O}_X(D+D_a-e_1-\cdots-e_m)\to\mathcal{O}_X(D+D_a)\to\bigoplus_i\mathcal{O}_{e_i}(D\cdot e_i)\to 0 \tag{10.72}$$

gives a complex

$$\mathcal{O}_X(e_1+\cdots+e_m-D_a+D)\to\mathcal{O}_X(D)^2\to\mathcal{O}_X(D+D_a) \tag{10.73}$$

representing $\bigoplus_i\mathcal{O}_{e_i}(D\cdot e_i)$. Taking the direct image in the case $D=(l-1)D_a$ gives a complex

$$\mathcal{O}_X((l-2)D_a)\to\mathcal{O}_X((l-1)D_a)^2\to\mathcal{O}_X(lD_a) \tag{10.74}$$

representing $\oplus_i\mathcal{O}_{qy_i}$ for each l. Since $R\,\mathrm{Hom}(\mathcal{O}_X(l_1D_a),\mathcal{O}_X(l_2D_a))$ is concentrated in degree 0 for $l_1\le l_2+1$, we find that the identity on this sum of point sheaves induces a chain map

$$\begin{array}{ccccc}
\mathcal{O}_X((l-3)D_a) & \longrightarrow & \mathcal{O}_X((l-2)D_a)^2 & \longrightarrow & \mathcal{O}_X((l-1)D_a)\\
\downarrow & & \downarrow & & \downarrow\\
\mathcal{O}_X((l-2)D_a) & \longrightarrow & \mathcal{O}_X((l-1)D_a)^2 & \longrightarrow & \mathcal{O}_X(lD_a)
\end{array} \tag{10.75}$$

the cone of which is a linear complex which is acyclic (and so certainly represents a sheaf). □

The analogous argument on X_{-1} gives the following.

Corollary 10.34. *The divisors h and $2h$ on X_{-1} satisfy the hypothesis of Lemma 10.32.*

Remark. We will see in Appendix B below that these cases actually *do* correspond to embeddings (in at least a weak sense) in noncommutative $\mathbb{P}^n$s.

Corollary 10.35. *If $m\le 5$ and C_m is ample, then $D_a=C_m$ satisfies the hypotheses of Lemma 10.32.*

Proof. If we restrict the four-term acyclic complex from Corollary 10.33 or Corollary 10.34 to C, we obtain (up to twisting by a line bundle) an acyclic complex of the form

$$\mathcal{L}_3\to\mathcal{L}_2^3\to\mathcal{L}_1^3\to\mathcal{O}_C \tag{10.76}$$

where $\mathcal{L}_i$ are line bundles such that $\mathcal{L}_3\otimes\mathcal{L}_1^3\cong\mathcal{L}_2$. Moreover, by varying the surface we use, we may arrange for $\mathcal{L}_1$ and $\mathcal{L}_2$ to be *arbitrary* bundles of the appropriate degrees.

In particular, we deduce that there is an acyclic complex of the form

$$\mathcal{O}_X(-3C_m)|_C \to \mathcal{O}_X(-2C_m)|_C^3 \to \mathcal{O}_X(-C_m)|_C^3 \to \mathcal{O}_C \tag{10.77}$$

on our original choice of X_m. If we represent each line bundle on C via a complex $\mathcal{O}_X(-(l+1)C_m) \to \mathcal{O}_X(-lC_m)$, then all of the relevant Ext groups vanish, and thus the above complex on C induces an acyclic complex of the form

$$\mathcal{O}_X(-4C_m) \to \mathcal{O}_X(-3C_m)^4 \to \mathcal{O}_X(-2C_m)^6 \to \mathcal{O}_X(-C_m)^4 \to \mathcal{O}_X. \tag{10.78}$$

Applying powers of θ gives a family of sheaves with linear resolutions as required. □

We do not know how to apply this to general very ample divisor classes, but the above is enough to let us prove the following.

Proposition 10.36. *Suppose M is a flat family of coherent sheaves on X_m over a noetherian base. Then the set of divisors D such that every fiber of $M(D)$ is acyclic and globally generated contains a translate of the nef cone.*

Proof. We first show that there is *some* divisor class such that $M(D)$ is acyclic and globally generated. To show this, first note that for any fiber M_v, if $\mathrm{Hom}(M_v, \mathcal{O}_{e_m}(-2)) = 0$ then $R^1\alpha_{m*}M_v = 0$; if M_v is pointless this follows from Lemma 11.17, and in general it follows by writing M_v as an extension of a pointless sheaf by a 0-dimensional sheaf. Now, $\mathrm{Hom}(M_v, \mathcal{O}_{e_m}(-2)) = 0$ iff $\mathrm{Ext}^2(\mathcal{O}_{e_m}(-1), M_v) = 0$, and this is an open condition by semicontinuity. Moreover, since $\theta\mathcal{O}_{e_m}(-1) \subset \mathcal{O}_{e_m}(-1)$, we find that $\mathrm{Ext}^2(\mathcal{O}_{e_m}(-1), M_v) = 0$ implies $\mathrm{Ext}^2(\mathcal{O}_{e_m}(-1), \theta^{-1}M_v)$. We thus obtain a nested sequence of closed subschemes of the base by taking those points where the fiber satisfies $\mathrm{Ext}^2(\mathcal{O}_{e_m}(-1), \theta^{-r}M_v) \neq 0$. Since $R^1\alpha_{m*}\theta^{-r}M_v$ is eventually acyclic, this nested sequence has empty intersection, so by the noetherian hypothesis is eventually empty. In particular, there exists r such that for *every* fiber, $R^1\alpha_{m*}\theta^{-r}M_v = 0$. If $(\alpha_{m*}\theta^{-r}M_v)(D)$ is globally generated, then so is $M_v(D + rC_m)$, so that the claim reduces by induction to the case $m \leq 0$.

For $m \leq 0$, we choose D_a as in Corollary 10.33 or Corollary 10.34 and take the nested sequence of closed subschemes of the base to be those fibers for which M_v is not (l, D_a)-regular.

We have thus shown that there is a divisor class D_1 such that for every fiber of the family, $\mathrm{Ext}^p(\mathcal{O}_X(-D_1), M_v) = 0$ for $p > 0$ and

$$\mathcal{O}_X(-D_1) \otimes_k \mathrm{Hom}(\mathcal{O}_X(-D_1), M_v) \to M_v \tag{10.79}$$

is surjective. Acyclicity implies that $\mathcal{H}om(\mathcal{O}_X(-D_1), M)$ is a flat coherent sheaf on the base, and thus the kernel of this family of surjective maps is itself a flat family over a noetherian base. We thus conclude that there is an exact sequence of the form

$$0 \to M' \to \mathcal{O}_X(-D_2) \otimes_k V \to \mathcal{O}_X(-D_1) \otimes_k \mathcal{H}om(\mathcal{O}_X(-D_1), M) \to M \to 0, \tag{10.80}$$

where again V is a flat coherent sheaf on the base and M' is a flat family of coherent sheaves on X_m. Now, if $D - D_1$ and $D - D_2$ are both very ample, then $\mathcal{O}_X(D - D_1)$ and $\mathcal{O}_X(D - D_2)$ are acyclic and globally generated. Since $M(D)$ is then a quotient of a globally generated sheaf, it is globally generated as required.

To see that $M(D)$ is acyclic, we use the hyper-Ext spectral sequence, and observe that the terms coming from $R\Gamma(M'(D))$ cannot cancel the terms corresponding to $\mathrm{Ext}^1(M(D))$ and $\mathrm{Ext}^2(M(D))$, which must therefore vanish as required. □

Remark. Note that we can allow the surface to vary in the family, as long as we work relative to the intersection of the nef cones of the fibers.

We will need a further boundedness result, but this depends on the theory of Hilbert schemes of points, so will be proved as Lemma 11.44 below.

Theorem 10.37. *Let M be a coherent sheaf on $X_{\rho;q;C}$, and let* $\mathrm{Quot}(M;(r,D,l))$ *denote the moduli functor classifying quotients of M with numerical invariants (r,D,l). Then* $\mathrm{Quot}(M;(r,D,l))$ *is represented by a projective scheme.*

Proof. Since the claim is clearly invariant under twisting, we may as well assume that M is acyclic and globally generated. But then $\mathrm{Quot}(M;(r,D,l))$ is a closed subfunctor of $\mathrm{Quot}(\mathcal{O}_X \otimes_k \mathrm{Hom}(\mathcal{O}_X, M);(r,D,l))$, so that it suffices to consider $M = \mathcal{O}_X^n$. Let D_a be a universally very ample divisor class. By Lemma 11.44 below, there is a bound B such that for all short exact sequences

$$0 \to I \to \mathcal{O}_X^n \to N \to 0 \tag{10.81}$$

corresponding to points of the Quot functor, $I(bD_a)$ and $N(bD_a)$ are acyclic and globally generated for $b \geq B$.

Now, let M_0 be the module associated to $\mathcal{O}_X$ over the $\mathbb{Z}$-algebra corresponding to D_a, and let M_B be the truncation obtained by removing the subspaces of degree $< B$. Then by the straightforward extension of [5, Thm. E4.3] to $\mathbb{Z}$-algebras, the quotients of M_B^n that have dimension $\chi(N(bD_a))$ for all $b \geq B$ are classified by a projective scheme. Let

$$Q^+ = 0 \to I^+ \to \mathcal{O}_X \to N^+ \to 0 \tag{10.82}$$

be the induced family of short exact sequences of sheaves on X. By our choice of B, this contains a universal family for the Quot scheme, and thus it remains only to show that the condition that a given fiber of N^+ has the desired numerical invariants is both closed and open. But this follows from the fact that the function $\chi(N^+(D))$ is locally constant for any D. □

Remark. Since we took D_a to be universally very ample, the argument actually works to construct the relative Quot scheme over any noetherian base, with both M and the surface allowed to vary. The only technical issue is that we can uniformly twist M to be acyclic and globally generated, but this will be shown in Proposition 10.36 below.

Remark. Traditionally, this is stated in terms of Hilbert polynomials, but for our purposes this is less natural. However, the same argument used to prove Lemma 11.44 shows that B can be chosen to depend only on the Hilbert polynomial, in which case N^+ represents the moduli functor classifying quotients with the chosen Hilbert polynomial.

Putting the above results together gives the following.

Theorem 10.38. *The quasi-scheme $X_{\rho;q;C}$ is a noncommutative smooth proper surface in the sense of [14].*

Proof. All of the axioms are either proved explicitly above or follow using the same arguments as in [14] for the ruled surface case. □

11 Moduli spaces

11.1 Poisson structures

Now that we know $X_{\rho;q;C}$ is a well-behaved noncommutative surface, our next objective is to understand moduli spaces of sheaves on $X_{\rho;q;C}$. In particular, since noncommutative surfaces are deformations of Poisson surfaces, and moduli spaces of sheaves on Poisson surfaces are Poisson [63, 12, 25], we would like to prove something similar in this setting. In addition, it was shown in [45] for the commutative setting that there are a number of natural Poisson morphisms between moduli spaces arising from birational morphisms, which again we would like to extend.

The first thing we need to do is construct the Poisson moduli space itself. The discussions in [63, 12, 25] concern moduli spaces of *stable* sheaves, which appear to be difficult to construct in general (see the next section for some special cases), as they require boundedness results. However, as observed in [45], a discussion of Poisson structures makes sense on general algebraic spaces, and there is a general construction [1] of an algebraic space representing a moduli problem that includes all stable sheaves. Thus our first order of business is to extend this construction to the noncommutative setting.

If I_S is a flat family of coherent sheaves on $X_{\rho;q;C}$ (where as usual (ρ, q, C) may vary over the base scheme S), we follow [1] in calling I_S *simple* if for any base change, the natural map

$$\mathcal{O}_T \to \mathcal{H}om_T(I_T, I_T) \tag{11.1}$$

is an isomorphism. Note that Serre duality allows us to rephrase the condition as saying that

$$\mathcal{H}om_T(\mathcal{E}xt^2_T(I_T, I_T), \mathcal{O}_T) \to \mathcal{O}_T \tag{11.2}$$

is an isomorphism. This makes the proof of [1, Prop. 5.2] carry through, so that "simple" is an open condition. (To be precise, there is an open, retrocompact subscheme of the base such that a morphism factors through the subscheme iff the base change is simple.) This allows us to rephrase the condition as saying that every endomorphism of a geometric fiber of I_S is a multiple of the identity. In addition, this allows us to define a functor $\mathcal{S}pl_{X_{\rho;q;C}/S}$ such that for any S-scheme T, $\mathcal{S}pl_{X_{\rho;q;C}/S}(T)$ is the set of equivalence classes of simple families over T, where two families are equivalent if there is an isomorphism

$$I \otimes_S L \cong J \tag{11.3}$$

where L is an invertible sheaf on S. Similarly, define a *twisted family of simple sheaves* over T to be a family of simple sheaves over an étale cover $U \to T$ such that the two induced families on $U \times_T U$ are equivalent; two such families are equivalent if they become equivalent under base change to a common étale cover.

Theorem 11.1. *There is a quasi-separated algebraic space $\mathcal{S}pl_{X_{\rho;q;C}/S}$ locally finitely presented over S which represents the moduli functor of simple sheaves, in the sense that there is a natural bijection between twisted families of simple sheaves on $X_{\rho;q;C}$ and morphisms to $\mathcal{S}pl_{X_{\rho;q;C}/S}$.*

Proof. We adapt the argument of [1, Thm. 7.4], with some minor modifications to avoid using Castelnuovo-Mumford regularity. Fix a universally ample divisor class D. For each pair r, n of positive integers, we may consider the open subfunctor $\Sigma_{r,n}$ classifying (twisted families of) simple sheaves M such that $\mathrm{Ext}^i(\mathcal{O}_X(-rD), M) = 0$ for $i > 0$, $\chi(\mathcal{O}_X(-rD), M) = n$, and the natural morphism $\mathcal{O}_X(-rD) \otimes \mathrm{Hom}(\mathcal{O}_X(-rD), M) \to M$ is surjective. As r and n vary, these subfunctors certainly cover $\mathcal{S}pl_{X_{\rho;q;C}/S}$, so it suffices to show that these subfunctors are represented

by algebraic spaces. Now, consider the projective scheme $\mathrm{Quot}(\mathcal{O}_X(-rD)^n)$.[3] This scheme has a universal family, and thus there is an open subscheme parametrizing simple sheaves such that $\mathrm{Ext}^i(\mathcal{O}_X(-rD), M) = 0$ for $i > 0$ and $\chi(\mathcal{O}_X(-rD), M) = n$; we may furthermore impose the open condition that the induced map

$$\mathrm{Hom}(\mathcal{O}_X(-rD), \mathcal{O}_X(-rD)^n) \to \mathrm{Hom}(\mathcal{O}_X(-rD), M) \tag{11.4}$$

is an isomorphism. This open subscheme $Z_{r,n}$ still covers $\Sigma_{r,n}$, and is in fact a PGL_n-bundle over $\Sigma_{r,n}$: it represents the moduli problem in which we include the additional data of a basis of $\mathrm{Hom}(\mathcal{O}_X(-rD), M)$, with proportional bases being equivalent. Moreover, given any morphism to $\Sigma_{r,n}$, the pullback to $Z_{r,n}$ is quasi-projective, as it is an open subscheme of

$$\mathbb{P}(\mathrm{Ext}^2(M, \theta\mathcal{O}_X(-rD)^n)); \tag{11.5}$$

the sheaf $\mathrm{Ext}^2(M, \theta\mathcal{O}_X(-rD)^n)$ is locally free since it is the only nonvanishing Ext sheaf. It follows as in [1] that $\Sigma_{r,n}$ is representable by the quotient $Z_{r,n}/\mathrm{PGL}_n$, and the claim follows. □

The covering by quotients of Quot schemes allows us to compute the cotangent sheaf, as in [32, 45]. Note that an étale morphism $T \to \mathcal{S}pl_{X_{rho;q;C}/S}$ is a formally universal *twisted* family of simple sheaves, but by definition there is always an étale cover $U \to T$ such that the twisted family becomes an honest family. In particular, any étale cover of $\mathcal{S}pl_{X_{\rho;q;C}/S}$ can be refined to a cover by formally universal families, and thus sheaves on $\mathcal{S}pl_{X_{\rho;q;C}/S}$ can be described in terms of the induced sheaves on the base of such families.

Lemma 11.2. *Let M be a formally universal family of simple sheaves on $X_{\rho;q;C}/S$ with base scheme U. Then there is a natural isomorphism*

$$\Omega_U \cong \mathcal{E}xt^1_U(M, \theta M). \tag{11.6}$$

This allows one to define a biderivation (following [63]) on $\mathcal{S}pl_{X_{\rho;q;C}/S}$, or equivalently on the base of every formally universal family, by

$$\Omega_U^{\otimes 2} \cong \mathcal{E}xt^1_U(M, \theta M)^{\otimes 2} \to \mathcal{E}xt^1_U(M, M) \otimes \mathcal{E}xt^1_U(M, \theta M) \to \mathcal{O}_U. \tag{11.7}$$

Note that for this to make sense, we must fix a natural transformation $\theta \to \mathrm{id}$. In particular, if we fix a nonzero holomorphic differential ω on C, then this allows us to identify θ and $\bar{\theta}$, at which point T gives such a natural transformation. (Such a choice also appears in the commutative setting as the difference between a choice of anticanonical curve and a choice of nontrivial Poisson structure.)

We then have the following result.

Theorem 11.3. *Let k be an algebraically closed field, let (C, ρ, q, ω) be a quadruple defined over k. Then the above biderivation equips $\mathcal{S}pl_{X_{\rho;q;C}}$ with a Poisson structure, and for any complex $M^{\cdot}_C$ of locally free sheaves on C, the subspace of $\mathcal{S}pl_{X_{\rho;q;C}}$ parametrizing sheaves M with $M|^{\mathbf{L}}_C \cong M^{\cdot}_C$ is a (smooth) symplectic Poisson subspace.*

Remark. As in [45], the claim that the biderivation gives a Poisson structure actually works fine in families, but the symplectic leaf claim has difficulties over more general schemes.

[3] Note that we only need that it is locally of finite type and a countable union of projective subschemes, which follows by the same argument as Theorem 10.37 without needing the boundedness result we have yet to prove.

The argument in [12, 25] for the commutative analogue of this fact (modulo the general claim about symplectic leaves) works by first showing that it holds for the subspace parametrizing vector bundles, and then showing that there is an open covering by subspaces isomorphic to open subspaces parametrizing vector bundles, in such a way that the isomorphism respects the pairing on the cotangent sheaf. Of course, there is a major difficulty in applying this argument in the noncommutative case: we do not have a natural notion of vector bundles on $X_{\rho;q;C}$ (and certainly cannot expect to do explicit calculations involving transition matrices!). That said, the reduction step itself works, and will prove useful.

Following [25] (see also the discussion in [45]), let $U \to \mathcal{S}pl_{X_{\rho;q;C}}$ be an étale neighborhood with corresponding family M of simple sheaves, where we suppose U noetherian. There thus exists a divisor class D such that $\mathcal{E}xt^i_U(\mathcal{O}_X(-D), M) = \mathcal{E}xt^i_U(\mathcal{O}_X(-D), \theta M) = 0$ for $i > 0$ and $\mathcal{O}_X(-D) \otimes_U \mathcal{E}xt^i_U(\mathcal{O}_X(-D), M) \to M$ is surjective. We may then consider the short exact sequence

$$0 \to V \to \mathcal{O}_X(-D) \otimes_U \mathcal{E}xt^i_U(\mathcal{O}_X(-D), M) \to M \to 0 \tag{11.8}$$

Then one can show (indeed, the argument as described in [45] carries over directly; "locally free" is only used to claim that one has reduced to the vector bundle case) that V is simple, and thus defines a new morphism $U \to \mathcal{S}pl_{X_{\rho;q;C}}$. Moreover, one can reconstruct M from V, and the two induced biderivations on U are the same (up to an overall sign). Finally, the derived restrictions $M|^{\mathbf{L}}_C$ and $V|^{\mathbf{L}}_C$ determine each other (once we fix the class of M in $K_0(X_{\rho;q;C})$). In other words, the Theorem holds on the étale neighborhood corresponding to M iff it holds on the étale neighborhood corresponding to V.

We should note here that this reduction applies equally well when M has a 0-dimensional subsheaf (at which point simplicity implies that M is 0-dimensional); in the commutative case, this was unnecessary since the only simple 0-dimensional sheaves are structure sheaves of points, but in the noncommutative setting, there are examples of any Euler characteristic $\leq |\langle q \rangle|$, see below.

This reduction, though it does not lead to a relatively straightforward case like vector bundles, does improve the sheaves under consideration in some respects. First, since the middle of the short exact sequence defining V is a sum of line bundles, so is acyclic for $|^{\mathbf{L}}_C$, V must also be acyclic for $|^{\mathbf{L}}_C$ (since derived restriction has homological dimension 1). If M were already acyclic for $|^{\mathbf{L}}_C$, then $V|_C$ would be a subsheaf of some power $\mathcal{O}_X(-D)^N|_C$, and thus torsion-free. And in either case, since $V \subset \mathcal{O}_X(-D)^N$, it must be a pure 2-dimensional sheaf. Finally, if we twist by D (which has no effect on the biderivation), we find that $R\Gamma(V) = 0$.

At this point, we make use of the following (rather surprising) fact. Let $\mathcal{N}_{\rho;q;C}$ denote the triangulated subcategory of $D^b \operatorname{coh} X_{\rho;q;C}$ consisting of objects M such that $R\Gamma M = 0$.

Proposition 11.4. *The triangulated subcategory $\mathcal{N}_{\rho;q;C}$ is independent of q. More precisely, for any q, there is an equivalence $\kappa_q : \mathcal{N}_{\rho;q;C} \cong \mathcal{N}_{\rho;1;C}$ such that*

$$(\kappa_q M)|^{\mathbf{L}}_C \cong M|^{\mathbf{L}}_C \otimes q. \tag{11.9}$$

Proof. We have already seen that $D^b \operatorname{coh} X_{\rho;q;C}$ is generated by the exceptional collection

$$\mathcal{O}_X(e_m)/\mathcal{O}_X, \mathcal{O}_X(e_{m-1})/\mathcal{O}_X, \ldots \mathcal{O}_X(e_1)/\mathcal{O}_X, \mathcal{O}_X(-s-2f), \mathcal{O}_X(-s-f), \mathcal{O}_X(-f), \mathcal{O}_X. \tag{11.10}$$

Moreover, $D^b \operatorname{coh} X_{\rho;q;C}$, as a derived category, comes with an enhancement to a dg-category (see, e.g., [40]), and that enhancement is equivalent to the (subcategory of compact objects of the) dg-category of modules over the exceptional collection. The same is therefore true for $\mathcal{N}_{\rho;q;C}$, where we simply omit $\mathcal{O}_X$ from the exceptional collection. Now, in general, if M and N are two objects with $R\operatorname{Hom}(N, M) = 0$, then $R\operatorname{Hom}(M, \theta N) = 0$ by Serre duality, and thus

$$R\operatorname{Hom}(M, N) \cong R\operatorname{Hom}(M, N|^{\mathbf{L}}_C) \cong R\operatorname{Hom}_C(M|^{\mathbf{L}}_C, N|^{\mathbf{L}}_C). \tag{11.11}$$

But this means that the dg-subcategory with objects the exceptional collection can be obtained from the corresponding subcategory of $D^b \operatorname{coh} C$ by simply omitting all backwards morphisms and nonidentity endomorphisms.

The restrictions of the objects in the exceptional collection can be computed from the fact that they are sheaves with known class in $K_0(C)$, giving

$$\mathcal{O}_{x_m}, \mathcal{O}_{x_{m-1}}, \dots, \mathcal{O}_{x_1}, \rho(-s-2f) \otimes q^{-1}, \rho(-s-f) \otimes q^{-1}, \rho(-f) \otimes q^{-1}. \tag{11.12}$$

If we tensor this with q, the result is independent of q, and the claim follows immediately. □

Remark. Note that the derived category $D^b \operatorname{coh} X_{\rho;q;C}$ can be reconstructed from the sequence

$$\mathcal{O}_{x_m}, \mathcal{O}_{x_{m-1}}, \dots, \mathcal{O}_{x_1}, \rho(-s-2f), \rho(-s-f), \rho(-f), q; \tag{11.13}$$

we will use this below to construct derived equivalences.

In principle, κ_q depends strongly on the particular choice of exceptional collection used to generate $\mathcal{N}_{\rho;q;C}$. Luckily, this dependence is relatively mild, and indeed, κ_q commutes with the action of admissible elements of $W(E_{m+1})$. This follows in a straightforward way from the following Lemma.

Lemma 11.5. *For all $0 \le l \le m$, and all integers d,*

$$\kappa_q(\mathcal{O}_X(-s-df+e_1+\cdots+e_l)) \cong \mathcal{O}_X(-s-df+e_1+\cdots+e_l). \tag{11.14}$$

In addition, $\kappa_q(\mathcal{O}_X(-2s-2f+e_1+\cdots+e_l)) \cong \mathcal{O}_X(-2s-2f+e_1+\cdots+e_l)$, $\kappa_q(\mathcal{O}_X(s)/\mathcal{O}_X) \cong \mathcal{O}_X(s)/\mathcal{O}_X$, and $\kappa_q(\mathcal{O}_X(f-e_1)/\mathcal{O}_X) \cong \mathcal{O}_X(f-e_1)/\mathcal{O}_X$.

Proof. For the first claim for $l=0$, we note the short exact sequences

$$0 \to \mathcal{O}_X(-s-(d-1)f) \to \mathcal{O}_X(-s-df)^2 \to \mathcal{O}_X(-s-(d+1)f) \to 0 \tag{11.15}$$

for all d. Since the claim holds for $d \in \{1,2\}$ and the morphisms in the sequence are canonical, the claim holds for all d. Similarly, we have a canonical short exact sequence

$$0 \to \mathcal{O}_X(-s-df+e_1+\cdots+e_{l-1}) \to \mathcal{O}_X(-s-df+e_1+\cdots+e_l) \to \mathcal{O}_X(e_l)/\mathcal{O}_X \to 0, \tag{11.16}$$

and thus the claim holds for all l by induction.

For $-2s-2f$, consider the short exact sequence

$$0 \to \mathcal{O}_X(-2s-2f) \to \mathcal{O}_X(-s-f) \oplus \mathcal{O}_X(-s-2f) \to \mathcal{O}_X(-f) \to 0. \tag{11.17}$$

The map $\mathcal{O}_X(-s-f) \oplus \mathcal{O}_X(-s-2f) \to \mathcal{O}_X(-f)$ is in the unique dense orbit of such maps under the action of $\operatorname{Aut}(\mathcal{O}_X(-s-f) \oplus \mathcal{O}_X(-s-2f))$, and thus this will remain true after applying κ_q. The extension to all l is as above.

Finally, we note that we have short exact sequences

$$0 \to \mathcal{O}_X(-2s-2f) \to \mathcal{O}_X(-s-2f) \to \mathcal{O}_X(s)/\mathcal{O}_X \to 0 \tag{11.18}$$

and

$$0 \to \mathcal{O}_X(-s-2f+e_1) \to \mathcal{O}_X(-s-f) \to \mathcal{O}_X(f-e_1)/\mathcal{O}_X \to 0, \tag{11.19}$$

with the first morphism begin unique in each case; this gives the final claims. □

Proposition 11.6. *If w is an admissible element of $W(E_{m+1})$ on $X_{\rho;q;C}$ (so also on $X_{\rho;1;C}$), then $\kappa_q \circ w = w \circ \kappa_q$.*

Proof. We observe that the Lemma allows us to compute κ_q using the exceptional collections

$$\mathcal{O}_X(e_m)/\mathcal{O}_X, \mathcal{O}_X(e_{m-1})/\mathcal{O}_X, \dots \mathcal{O}_X(e_1)/\mathcal{O}_X, \mathcal{O}_X(s)/\mathcal{O}_X, \mathcal{O}_X(-2s-2f), \mathcal{O}_X(-s-f) \quad (11.20)$$

and

$$\mathcal{O}_X(e_m)/\mathcal{O}_X, \mathcal{O}_X(e_{m-1})/\mathcal{O}_X, \dots \mathcal{O}_X(f-e_1)/\mathcal{O}_X, \mathcal{O}_X(-s-2f+e_1), \mathcal{O}_X(-s-f+e_1), \mathcal{O}_X(-f). \quad (11.21)$$

Each admissible simple reflection of $W(E_{m+1})$ acts by swapping two adjacent orthogonal elements of one of these two exceptional collections, and thus commutes with κ_q. □

Note that $R\alpha_{m*}, L\alpha_m^*$ preserve the condition that $R\Gamma = 0$, and thus induce functors between the respective categories $\mathcal{N}_{\rho;q;C}$.

Corollary 11.7. *One has natural isomorphisms*

$$\kappa_q R\alpha_{m*} \cong R\alpha_{m*}\kappa_q,$$
$$\kappa_q L\alpha_m^* \cong L\alpha_m^*\kappa_q,$$
$$\kappa_q L\alpha_m^! \cong L\alpha_m^!\kappa_q,$$

as well as

$$\kappa_q R\operatorname{ad} \cong R\operatorname{ad}\kappa_q. \quad (11.22)$$

Proof. Indeed, in the first two cases, both functors have the same action on the elements of the appropriate exceptional collection. (For $m = 0$, use the $\mathbb{P}^2$-based exceptional collection instead.) For the third case, use the distinguished triangles

$$R\operatorname{Hom}(\mathcal{O}_{e_m}(-1), L\alpha_m^*\kappa_q M) \otimes_k^{\mathbf{L}} \mathcal{O}_{e_m}(-1) \to L\alpha_m^*\kappa_q M \to L\alpha_m^!\kappa_q M \to \quad (11.23)$$

and

$$R\operatorname{Hom}(\mathcal{O}_{e_m}(-1), L\alpha_m^* M) \otimes_k^{\mathbf{L}} \mathcal{O}_{e_m}(-1) \to \kappa_q L\alpha_m^* M \to \kappa_q L\alpha_m^! M \to \quad (11.24)$$

together with the fact that

$$R\operatorname{Hom}(\mathcal{O}_{e_m}(-1), L\alpha_m^* M) \cong R\operatorname{Hom}(\mathcal{O}_{e_m}(-1), \kappa_q L\alpha_m^* M) \quad (11.25)$$

to obtain compatible isomorphisms between the first two terms, making the third terms isomorphic.

Finally, for $R\operatorname{ad}$, we note that the Lemma directly tells us that κ_q and $R\operatorname{ad}$ commute on the line bundles in the exceptional collection. For the sheaves $\mathcal{O}_X(e_l)/\mathcal{O}_X$, we have

$$\mathcal{O}_X(e_l)/\mathcal{O}_X \cong L\alpha_m^* \cdots L\alpha_{l+1}^* \mathcal{O}_{e_l}(-1), \quad (11.26)$$

and thus

$$R\operatorname{ad}(\mathcal{O}_X(e_l)/\mathcal{O}_X) \cong L\alpha_m^! \cdots L\alpha_{l+1}^! \mathcal{O}_{e_l}(-1), \quad (11.27)$$

and this is preserved by κ_q. □

Since $\mathcal{N}_{\rho;q;C}$ is compactly generated, there is a right adjoint $\Phi_{\rho;q;C}$ to the functor $|_C^{\mathbf{L}}$, and we immediately have the following.

Lemma 11.8. *We have a natural isomorphism* $\kappa_q \Phi_{\rho;q;C} M \cong \Phi_{\rho;1;C}(M \otimes q)$.

Now, define a functor $\bar{\theta}_{\mathcal{N}} : \mathcal{N}_{\rho;q;C} \to D^b \operatorname{coh} X_{\rho;q;C}$ via the distinguished triangle

$$\bar{\theta}_{\mathcal{N}} M \to M \to \Phi_{\rho;q;C} M|_C^{\mathbf{L}} \to \tag{11.28}$$

and let $T_{\mathcal{N}}$ denote the natural transformation $\bar{\theta}_{\mathcal{N}} \to \mathrm{id}$. (Here we are using the dg enhancement to allow us to define the cone functor.) Let us also define $\theta_{\mathcal{N}} \cong \bar{\theta}_{\mathcal{N}} \otimes_k H^0_C(\omega_C)$.

Proposition 11.9. *We have* $\bar{\theta}_{\mathcal{N}} \kappa_q \cong \kappa_q \bar{\theta}_{\mathcal{N}}$ *and* $T_{\mathcal{N}} \kappa_q \cong \kappa_q T_{\mathcal{N}}$. *Moreover,* $\mathcal{N}_{\rho;q;C}$ *satisfies Serre duality in the form of functorial isomorphisms*

$$\operatorname{Ext}^i(M, N) \cong \operatorname{Ext}^{2-i}(N, \theta_{\mathcal{N}} M). \tag{11.29}$$

Proof. The first two claims follow immediately from the Lemma. For the remaining claim, note that we have a distinguished triangle

$$R\operatorname{Hom}(N, \bar{\theta}_{\mathcal{N}} M) \to R\operatorname{Hom}(N, M) \to R\operatorname{Hom}(N, \Phi_{\rho;q;C} M|_C^{\mathbf{L}}) \to \tag{11.30}$$

which by adjunction is naturally equivalent to

$$R\operatorname{Hom}(N, \bar{\theta}_{\mathcal{N}} M) \to R\operatorname{Hom}(N, M) \to R\operatorname{Hom}_C(N|_C^{\mathbf{L}}, M|_C^{\mathbf{L}}) \to \tag{11.31}$$

and thus to

$$R\operatorname{Hom}(N, \bar{\theta}_{\mathcal{N}} M) \to R\operatorname{Hom}(N, M) \to R\operatorname{Hom}(N, M|_C^{\mathbf{L}}) \to \tag{11.32}$$

We thus have a natural isomorphism

$$R\operatorname{Hom}(N, \bar{\theta}_{\mathcal{N}} M) \cong R\operatorname{Hom}(N, \bar{\theta} M) \tag{11.33}$$

and thus the claim follows from Serre duality on $X_{\rho;q;C}$. □

We may now finish the proof of Theorem 11.3. If M is an object in $\mathcal{N}_{\rho;q;C} \cap \operatorname{coh} X_{\rho;q;C}$ such that $\kappa_q M \in \operatorname{coh} X_{\rho;1;C}$, or more precisely a formally universal flat family of such objects, then we may compute the corresponding pairing on the cotangent sheaf in either of $X_{\rho;q;C}$ or $X_{\rho;1;C}$. But the two computations are the same: we are simply replacing θ by $\theta_{\mathcal{N}}$ in the definition of the biderivation (which has no effect on the Ext^1) and using Serre duality on $\mathcal{N}_{\rho;q;C}$ to define the pairing.

Thus to prove that the biderivation gives a Poisson structure, we simply need to show that we can arrange for our object $V \in \mathcal{N}_{\rho;q;C} \cap \operatorname{coh} X_{\rho;q;C}$ to correspond to a sheaf in $X_{\rho;1;C}$.

Lemma 11.10. *Let* M *be a sheaf on* $\operatorname{coh} X_{\rho;q;C}$, *and define a sequence of objects* $M_0, \ldots, M_m$ *by* $M_m = M$, $M_{l-1} = R\alpha_{l*} M_l$. *Suppose that*

(1) $R\operatorname{Hom}(\mathcal{O}_X, M)$, $R\operatorname{Hom}(\mathcal{O}_X(f), M)$, $R\operatorname{Hom}(\mathcal{O}_X(s+f), M)$, *and* $R\operatorname{Hom}(\mathcal{O}_X(s+2f), M)$ *are supported in degree 0.*

(2) Each M_l *is an* α_l^**-acyclic sheaf, and the morphisms* $L\alpha_l^* M_{l-1} \to M_l$ *are surjective.*

Finally, let $V \in \mathcal{N}_{\rho;q;C}$ *be the kernel in the short exact sequence*

$$0 \to V \to \Gamma(M) \otimes \mathcal{O}_X \to M \to 0. \tag{11.34}$$

Then $\kappa_q V$ *is a sheaf of homological dimension 1 on* $X_{\rho;1;C}$.

Proof. Let $V_m = V$, $V_{l-1} = R\alpha_{(l-1)*}V_l$. Repeatedly applying $R\alpha_*$ tells us that V_l is a sheaf for each $0 \leq l \leq m$ and fits into a short exact sequence

$$0 \to V_l \to \Gamma(M) \otimes \mathcal{O}_X \to M_l \to 0 \tag{11.35}$$

Suppose we know that $\kappa_q V_{l-1}$ is a sheaf of homological dimension 1. Then it is, in particular, α_l^*-acyclic, so that

$$\kappa_q L\alpha_l^* R\alpha_{l*} V_l = L\alpha_l^* \kappa_q R\alpha_{l*} V_l \tag{11.36}$$

is a sheaf. In particular, $L\alpha_l^* R\alpha_{l*}$ is exact on the above short exact sequence, and we may apply the snake lemma to the resulting map of short exact sequences. Since $L\alpha_l^* R\alpha_{l*} \to \mathrm{id}$ is an isomorphism in the middle, we conclude that we have a short exact sequence of the form

$$0 \to L\alpha_l^* R\alpha_{l*} V_l \to V_l \to \mathcal{O}_{e_l}(-1)^n \to 0. \tag{11.37}$$

This is a short exact sequence of sheaves in $\mathcal{N}_{\rho;q;C}$, so we may apply κ_q to obtain a distinguished triangle of such sheaves; since the kernel and cokernel are sheaves, so is $\kappa_q V_l$. Moreover, both kernel and cokernel have homological dimension 1, and thus the same holds for $\kappa_q V_l$.

Thus, by induction, we reduce to considering the case $l = 0$. Now, M is represented in $D^b \operatorname{coh} X_{\eta,x_0;q;C}$ by a bounded complex in which each term is a sum of sheaves $\mathcal{O}_X$, $\mathcal{O}_X(-f)$, $\mathcal{O}_X(-s-f)$, $\mathcal{O}_X(-s-2f)$, since that exceptional collection generates the category and there are no extensions between the sheaves. We may further arrange for there to be no nonzero maps in the complex between line bundles of the same divisor class. Taking (derived) homomorphisms from $\mathcal{O}_X$, $\mathcal{O}_X(f)$, $\mathcal{O}_X(s)$, $\mathcal{O}_X(s+f)$ allows us to compute the dimensions; since M_0 is acyclic in each case, we find that M_0 has a presentation of the form

$$0 \to \mathcal{O}_X(-s-2f)^{n_1} \to \mathcal{O}_X(-s-f)^{n_2} \oplus \mathcal{O}_X(-f)^{n_3} \to \mathcal{O}_X^{n_4} \to M_0 \to 0. \tag{11.38}$$

This in turn implies that V_0 has a presentation

$$0 \to \mathcal{O}_X(-s-2f)^{n_1} \to \mathcal{O}_X(-s-f)^{n_2} \oplus \mathcal{O}_X(-f)^{n_3} \to V_0 \to 0. \tag{11.39}$$

Since $|_C^{\mathbf{L}}$ has homological dimension 1, the map $\theta V_0 \to V_0$ is injective, which implies that the presentation for V_0 restricts to a presentation for $V_0|_C$. In particular, we find that the map of sums of line bundles is injective on C. Applying κ_q gives a morphism of vector bundles which is still injective on C, and is thus injective on $X_{\eta,x_0;1;C}$. It follows that $\kappa_q V_0$ is a sheaf of homological dimension 1 as required. □

At this point, it is straightforward to complete the proof of Theorem 11.3. The only thing we need to observe is that for any 2-dimensional sheaf M, there is a divisor class D such that $M(D)$ satisfies the above conditions, as well as $H^i(\theta M(D)) = 0$ for $i > 0$. Indeed, each requirement holds for D sufficiently ample, and thus there is a nonempty set of D satisfying all of the conditions. □

We next would like to understand how the various Poisson morphisms of [45] extend to the present setting. It turns out that much of that extension is entirely straightforward, with only a few places that require different arguments or hypotheses. The most significant of the latter is the requirement that various sheaves have homological dimension 1. This, of course, makes no sense in our setting, since we do not have a well-behaved notion of "locally free". However, from [45, Prop. 4.2], we see that a sheaf on a commutative surface has homological dimension 1 iff it contains no 0-dimensional subsheaf. The latter notion, of course, extends immediately to the noncommutative setting; we thus call a sheaf "pointless" if it has no 0-dimensional subsheaf. Note that there are some instances in which a sheaf on $X_{\rho;q;C}$ is inarguably of homological dimension 1 by any reasonable definition; luckily, those sheaves are indeed pointless.

Lemma 11.11. *The cokernel of an injective morphism between two sums of line bundles is pointless.*

Proof. This follows immediately via the long exact sequence from the fact that for any 0-dimensional sheaf M and divisor class D, $\mathrm{Ext}^i(M, \mathcal{O}_X(D)) = 0$ for $i < 2$. □

Lemma 11.12. *A simple coherent sheaf is either pointless or* 0*-dimensional.*

Proof. Let M be a simple coherent sheaf which is neither pointless nor 0-dimensional, and let N be the sum of all 0-dimensional subsheaves of M, making N the maximal 0-dimensional subsheaf of M. Then M is a nontrivial extension of M/N by N, so that $\mathrm{Ext}^1(M/N, N) \neq 0$. Since N is 0-dimensional, we have $\chi(M/N, N) = \mathrm{rank}(M/N)\chi(N) \geq 0$, and may thus conclude that

$$\dim \mathrm{Hom}(M/N, N) + \dim \mathrm{Hom}(\theta^{-1}N, M/N) = \dim \mathrm{Hom}(M/N, N) + \dim \mathrm{Ext}^2(M/N, N) > 0. \tag{11.40}$$

A nonzero morphism $M/N \to N$ would contradict simplicity of M, while a nonzero morphism $\theta^{-1}N \to M/N$ would contradict maximality of N. □

One helpful fact to establish is that sheaves on -1-curves are rigid. More precisely, we have the following. For each integer d, define

$$\mathcal{O}_{e_m}(d) = \mathcal{O}_X(-de_m)/\mathcal{O}_X(-(d-1)e_m), \tag{11.41}$$

which is pointless by the Lemma.

Lemma 11.13. *If M is a pointless coherent sheaf with numerical invariants $(0, e_m, d)$, then*

$$M \cong \mathcal{O}_{e_m}(d-1). \tag{11.42}$$

Proof. Since e_m is not a sum of two nonzero effective divisors (indeed, there are ample divisors D such that $D \cdot e_m = 1$), any map between sheaves of Chern class e_m has either 0-dimensional image or 0-dimensional cokernel. In particular, if both sheaves are pointless, then the morphism must be injective, since neither the image nor the kernel can be 0-dimensional.

Since $\chi(M, \mathcal{O}_{e_m}(d-1)) = 1$, we conclude that at least one of the spaces $\mathrm{Hom}(M, \mathcal{O}_{e_m}(d-1))$ or $\mathrm{Ext}^2(M, \mathcal{O}_{e_m}(d-1)) \cong \mathrm{Hom}(\mathcal{O}_{e_m}(d-1), \theta M)^*$ is nonzero. In the first case, we have an injective morphism $M \to \mathcal{O}_{e_m}(d-1)$, and K_0 considerations tell us that the cokernel is trivial, and thus the map is an isomorphism. In the second case, the morphism must still be injective, but now the cokernel would be a 0-dimensional sheaf with Euler characteristic -1. □

Remark. In particular, $\theta\mathcal{O}_{e_m}(d) \cong \mathcal{O}_{e_m}(d-1)$, and any subsheaf of $\mathcal{O}_{e_m}(d)$ has the form $\mathcal{O}_{e_m}(d')$ for $d' < d$.

We may use this to give an example of exotic 0-dimensional simple sheaves.

Lemma 11.14. *If $d \leq |\langle q \rangle|$, then the cokernel of $T^d : \bar{\theta}^d\mathcal{O}_{e_m}(k) \to \mathcal{O}_{e_m}(k)$ is a simple sheaf for any k.*

Proof. Let $M_{d,k}$ be the given cokernel. It is clear that $M_{d,k}|_C \cong M_{1,k}$, and since T^d is injective for all d, we find $M_{1,k} \cong \mathcal{O}_{q^{k+1}x_m}$. Moreover, $M_{d,k}$ has a natural filtration with quotients $M_{1,k}$, $M_{1,k-1}, \ldots, M_{1,k-d+1}$.

Let d be minimal such that $M_{d,k}$ is not simple (note that since $\theta M_{d,k} \cong M_{d,k+1}$, simplicity is independent of k), and let ϕ be a nonscalar endomorphism of $M_{d,k}$. Consider the restriction of ϕ to the subsheaf $M_{d-1,k-1}$. If $\phi(M_{d-1,k-1}) \not\subset M_{d-1,k-1}$, then we can compose with the natural map

$M_{d,k} \to M_{d,k}|_C$ to obtain a map $M_{d-1,k-1} \to M_{1,k}$. Since $d \leq |\langle q \rangle|$, none of the subquotients of $M_{d-1,k-1}$ map to $M_{1,k}$, giving a contradiction.

It follows that $\phi(M_{d-1,k-1}) \subset M_{d-1,k-1}$. In that case, induction tells us that the restriction must be scalar, and thus subtracting that scalar gives a nonzero endomorphism that annihilates $M_{d-1,k-1}$, and thus a morphism $M_{1,k} \to M_{d,k}$. Since T annihilates the domain, it must also annihilate the image, and thus this morphism actually maps $M_{1,k} \to M_{1,k-d+1}$. But this cannot happen unless $d-1$ is a multiple of $|\langle q \rangle|$. □

Remark. Note that if $d = |\langle q \rangle|$, then there are other examples, coming from structure sheaves of points on the center.

We also need to know a bit about how these sheaves interact with $R\alpha_{m*}$, $L\alpha_m^*$, and $L\alpha_m^!$.

Lemma 11.15. *We have*

$$\begin{aligned} R\alpha_{m*}\mathcal{O}_{e_m} &\cong \mathcal{O}_{qx_m}, \\ R\alpha_{m*}\mathcal{O}_{e_m}(-1) &= 0 \\ R\alpha_{m*}\mathcal{O}_{e_m}(-2) &\cong \mathcal{O}_{x_m}[-1] \end{aligned}$$

In addition, for $y \in C$, $L\alpha_m^*\mathcal{O}_y \cong \mathcal{O}_y$ if $y \neq qx_m$, for which

$$\begin{aligned} \alpha_m^*\mathcal{O}_{qx_m} &\cong \mathcal{O}_{e_m} \\ L_1\alpha_m^*\mathcal{O}_{qx_m} &\cong \mathcal{O}_{e_m}(-1). \end{aligned}$$

Similarly, $L\alpha_m^!\mathcal{O}_y \cong \mathcal{O}_y$ for $y \neq x_m$, and

$$\begin{aligned} \alpha_m^!\mathcal{O}_{x_m} &\cong \mathcal{O}_{e_m}(-1) \\ L_1\alpha_m^!\mathcal{O}_{x_m} &\cong \mathcal{O}_{e_m}(-2). \end{aligned}$$

Proof. That $R\alpha_{m*}\mathcal{O}_{e_m}(-1) = 0$ was established in [64]. The other two claims for $R\alpha_{m*}$ then follow from the short exact sequences

$$\begin{aligned} 0 \to \mathcal{O}_{e_m}(-2) \to \mathcal{O}_{e_m}(-1) \to \mathcal{O}_{x_m} \to 0 \\ 0 \to \mathcal{O}_{e_m}(-1) \to \mathcal{O}_{e_m} \to \mathcal{O}_{qx_m} \to 0 \end{aligned}$$

The formula for $L\alpha_m^*\mathcal{O}_y$ is [64, Prop. 8.3.2], and the formula for $L\alpha_m^!$ is then straightforward. □

The discussion of [45] could be applied to any birational morphism, but since most of the arguments involved restricting to monoidal transformations, we will mainly consider that case; that is, how moduli spaces on X_m and X_{m-1} are related.

Lemma 11.16. *If M is a pointless coherent sheaf on X_{m-1}, then $L\alpha_m^*$ and $L\alpha_m^!$ are pointless α_{m*}-acyclic sheaves with direct image naturally isomorphic to M.*

Proof. We certainly have $M \cong R\alpha_{m*}L\alpha_m^*M$ by the natural map, and thus the natural spectral sequence gives:

$$\alpha_{m*}L_1\alpha_m^*M = R^1\alpha_{m*}\alpha_m^*M = 0 \tag{11.43}$$

along with a short exact sequence

$$0 \to R^1\alpha_{m*}L_1\alpha_m^*M \to M \to \alpha_{m*}\alpha_m^*M \to 0 \tag{11.44}$$

Since the image of $R^1\alpha_{m*}$ is always 0-dimensional ([64, Prop. 6.5.2]), pointlessness of M implies that $R\alpha_{m*}L_1\alpha_m^*M = 0$, so that $L_1\alpha_m^*M \cong \mathcal{O}_e(-1)^n$ for some n. But we then have, as above,

$$\mathrm{Hom}(\mathcal{O}_{e_m}(-1), L_1\alpha^*M) \cong \mathrm{Hom}(\mathcal{O}_{qx_m}, M) = 0 \tag{11.45}$$

since M is pointless. It follows that $L\alpha_m^*M$ is a sheaf with $\alpha_{m*}\alpha_m^*M \cong M$. The direct image of a 0-dimensional subsheaf of α_m^*M thus gives a 0-dimensional subsheaf of M, implying that $L\alpha_m^*M$ is pointless.

Then claim for $L\alpha_m^!M$ follows immediately, since θ takes pointless sheaves to pointless sheaves. □

Now, the natural isomorphism $R\alpha_{m*}L\alpha_m^* \to \mathrm{id}$ (the inverse of the natural map coming from the adjunction) is transformed by adjunction to a natural transformation $L\alpha_m^* \to L\alpha_m^!$.

Definition 11.1. Let M be a pointless coherent sheaf on X_{m-1}. Then the "minimal lift" of M is the image $\alpha_m^{*!}M$ of the morphism $\alpha_m^*M \to \alpha_m^!M$.

Remark. As in [45], this extends readily to a functor lifting pointless coherent sheaves from X_k to X_m, either by taking $\alpha_m^{*!}\alpha_{m-1}^{*!}\cdots\alpha_{k+1}^{*!}M$ or by taking the image of the natural map

$$\alpha_m^*\alpha_{m-1}^*\cdots\alpha_{k+1}^*M \to \alpha_m^!\alpha_{m-1}^!\cdots\alpha_{k+1}^!M. \tag{11.46}$$

The name "minimal lift" follows from the fact (with the same proof as in [45]) that any coherent α_{m*}-acyclic sheaf on X_m with direct image M has a subquotient isomorphic to $\alpha_m^{*!}M$, such that the other factors are isomorphic to powers of $\mathcal{O}_{e_m}(-1)$.

Much of the discussion in [45, §5] of the category of sheaves annihilated by $R\alpha_{m*}$ or more generally $R\alpha_{(k+1)*}\cdots R\alpha_{(m-1)*}R\alpha_{m*}$ carries over with little difficulty to the noncommutative setting. One notable exception is Proposition 5.4 op. cit., for which the hypotheses do not make sense (they refer to the locus where a morphism vanishes); this requires some modifications to later arguments. In Lemma 5.9 op. cit., "homological dimension ≤ 1" and "homological dimension 2" need to be replaced by "pointless" and "not pointless" respectively. There, we need to use the fact that a 0-dimensional subsheaf of the (higher) direct image of a pointless sheaf necessarily has the form $\mathcal{O}_p$ for $p \in C$. We record this here for future reference, including the analogue of Corollary 5.10 op. cit.

Lemma 11.17. *Let M be a pointless coherent sheaf on X_m. If M is not α_{m*}-acyclic, then $\mathrm{Hom}(M, \mathcal{O}_{e_m}(-2)) \neq 0$, while if $\alpha_{m*}M$ is not pointless, then $\mathrm{Hom}(\mathcal{O}_{e_m}, M) \neq 0$. Moreover, if $\mathrm{Hom}(\mathcal{O}_{e_m}(-1), M) = \mathrm{Hom}(M, \mathcal{O}_{e_m}(-1)) = 0$ then $R\alpha_{m*}M$ is a pointless sheaf, with $M \cong \alpha_m^{*!}\alpha_{m*}M$.*

In addition, there are natural subtleties arising around -2-curves; one finds that the simple objects in the kernel of $R\alpha_{(k+1)*}\cdots R\alpha_{(m-1)*}R\alpha_{m*}$ have the form $\mathcal{O}_X(e_i)/\mathcal{O}_X(e_j)$ if j is the minimal index such that $x_i = x_j$, and $\mathcal{O}_X(e_i)/\mathcal{O}_X$ if there is no such index. The key observation one must use to rule out other subsheaves is that a sheaf with Chern class $e_i - e_j$ has degree 0 derived restriction to C, and thus to be pointless must have *trivial* derived restriction. Since this is a subcategory of $\mathcal{N}_{\rho;q;C}$, we may use the functor κ_q to move everything to $\mathcal{N}_{\rho;1;C}$; this identifies the respective sets of simple objects, and thus identifies the actual abelian categories.

Continuing on with minimal lifts, one has the following, essentially combining the non-sheaf-Hom version of [45, Lem. 6.1] with [45, Prop. 6.4].

Lemma 11.18. *Suppose M and N are pointless coherent sheaves on X_{m-1}. Then there is an isomorphism*

$$\mathrm{Hom}(\alpha_m^{*!}M, \alpha_m^{*!}N) \cong \mathrm{Hom}(M, N) \tag{11.47}$$

and an exact sequence

$$\begin{aligned} 0 \to \mathrm{Ext}^1(\alpha_m^{*!}M, \alpha^{*!}N) \to \mathrm{Ext}^1(M, N) \\ \to \mathrm{Hom}_k(\mathrm{Hom}(N, \mathcal{O}_{x_m}), \mathrm{Ext}^1(M, \mathcal{O}_{x_m})) \to \mathrm{Ext}^2(\alpha_m^{*!}M, \alpha^{*!}N) \to 0. \end{aligned}$$

Proof. We have exact sequences

$$\begin{aligned} 0 \to E_1 \to \alpha_m^* M \to \alpha^{*!} M \to 0 \\ 0 \to \alpha^{*!} N \to \alpha^! N \to E_2 \to 0 \end{aligned}$$

where E_1 and E_2 are both powers of $\mathcal{O}_{e_m}(-1)$. We find

$$R\,\mathrm{Hom}(E_1, \alpha^{*!}N) \cong R\,\mathrm{Hom}(E_1, E_2)[-1] \cong R\,\mathrm{Hom}(\alpha^{*!}M, E_2), \tag{11.48}$$

and thus it remains only to determine E_1, E_2, or equivalently the spaces $\mathrm{Hom}(\mathcal{O}_{e_m}(-1), \alpha_m^* M)$ and $\mathrm{Hom}(\alpha_m^! M, \mathcal{O}_{e_m}(-1))$. For the first, we have

$$\mathrm{Hom}(\mathcal{O}_{e_m}(-1), \alpha_m^* M) \cong \mathrm{Hom}(\alpha_m^* M, \mathcal{O}_{e_m}(-2))^* \cong \mathrm{Ext}^1(M, \mathcal{O}_{x_m})^* \tag{11.49}$$

and a similar argument gives

$$\mathrm{Hom}(\alpha_m^! M, \mathcal{O}_{e_m}(-1)) \cong \mathrm{Hom}(M, \mathcal{O}_{x_m}). \tag{11.50}$$

□

Corollary 11.19. *Let M be a pointless coherent sheaf on X_{m-1}. Then the sheaves $\theta^{\pm 1}\alpha_m^{*!}M$ are α_{m*}-acyclic, with pointless direct image. Moreover, there are short exact sequences*

$$0 \to M \to \theta^{-1}\alpha_{m*}\theta\alpha_m^{*!}M \to \mathcal{O}_{qx_m} \otimes_k \mathrm{Ext}^1(\mathcal{O}_{qx_m}, M) \to 0 \tag{11.51}$$

and

$$0 \to \theta\alpha_{m*}\theta^{-1}\alpha_m^{*!}M \to M \to \mathrm{Hom}_k(\mathrm{Hom}(M, \mathcal{O}_{x_m}), \mathcal{O}_{x_m}) \to 0. \tag{11.52}$$

Remark. Note also that one can read off the dimensions of $\mathrm{Ext}^1(\mathcal{O}_{qx_m}, M)$ and $\mathrm{Hom}(M, \mathcal{O}_{x_m})$ from the numerical invariants of $\alpha_m^{*!}M$.

There is a subtlety in extending the proof of Theorem 6.6 of [45], as it involves comparing sheaves which are twists by different line bundles. We can replace this by an inductive argument as follows.

Lemma 11.20. *Let M, N be pointless sheaves on X_m and X_{m-1}. If there is a short exact sequence*

$$0 \to \alpha_m^{*!}N \to M \to \mathcal{O}_x^l \to 0 \tag{11.53}$$

for some $x \in C$, then $R\alpha_{m*}M$ is a pointless sheaf. Similarly, if there is a short exact sequence

$$0 \to M \to \alpha_m^{*!}N \to \mathcal{O}_x^l \to 0, \tag{11.54}$$

then $R\alpha_{m*}M$ is a pointless sheaf.

Proof. For the first claim, we note that $R\alpha_{m*}M$ is an extension of sheaves, so a sheaf, and it remains only to show that it is pointless. We may now observe that since $\mathrm{Hom}(\mathcal{O}_x, \mathcal{O}_{e_m}(-1)) = 0$, the pushforward exact sequence

$$0 \to \alpha_m^! N \to M' \to \mathcal{O}_x^l \to 0 \tag{11.55}$$

is non-split, and thus M' is an extension of a power of $\mathcal{O}_{e_m}(-1)$ by M, so remains pointless and has the same direct image as M. The claim then follows from the isomorphism

$$\mathrm{Ext}^1(\mathcal{O}_x, \alpha_m^! N) \cong \mathrm{Ext}^1(R\alpha_{m*}\mathcal{O}_x, N) \cong \mathrm{Ext}^1(\mathcal{O}_x, N), \tag{11.56}$$

which implies that the direct image of such a nonsplit extension remains nonsplit, making $\alpha_{m*}M = \alpha_{m*}M'$ pointless.

The other claim is analogous, where now the nontrivial task is to prove that $R\alpha_{m*}M$ is a sheaf. □

Given the Lemma, the analogue of [45, Thm. 6.6] is straightforward.

Theorem 11.21. *[45] Let M be a pointless coherent sheaf on X_k. Then the direct images in $D^b \operatorname{coh} X_k$ of $\theta^{\mp 1}\alpha_{m*}\theta^{\pm 1}\alpha_m^{*!}\cdots\alpha_{k+1}^{*!}M$ are pointless coherent sheaves.*

Another important observation is that (per [45, Lem. 6.8], with the same proof) given any pure 1-dimensional sheaf on X_m on which T is injective, we can always choose further blowups in such a way that the minimal lift to some X_{m+k} will have T an isomorphism. The argument lifts directly, with some simplification due to the fact that C is smooth, and thus "pseudo-twists" are not required. This is important for our interpretations and generalizations of Painlevé below; we will obtain Painlevé-like integrable systems via twisting morphisms between moduli spaces of pure 1-dimensional sheaves disjoint from C, and use the above results to see how the operations act in terms of difference equations.

In general, two difference equations are related by a gauge equivalence (a.k.a. isomonodromy deformation) iff their corresponding sheaves are *comparable*, in the sense that there is a common subsheaf with the same Chern class. Part of the justification of reducing to the "disjoint from C" case is that two equations are gauge equivalent iff their disjoint from C versions are equivalent via *canonical* gauge transformations (i.e., those coming from twists).

Proposition 11.22. *Suppose q is non-torsion and M, N are comparable pure 1-dimensional sheaves on X_m which are both transverse to C. Then M and N have a common pseudo-twist.*

Proof. We first observe that it suffices to consider the case $M \subset N$, as otherwise we may instead consider the pair $(M \cap N)^2 \subset M \oplus N$ where $M \cap N$ is the maximal common subsheaf. If M and N are disjoint from C, then their Euler characteristics must agree by Corollary 11.30 below, and thus $M = N$. Otherwise, choose a point x_{m+1} in the support of $N|_C$, and let X_{m+1} denote the blowup of X_m in that point. The snake lemma applied to

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \theta\alpha_{(m+1)*}\theta^{-1}\alpha_{m+1}^{*!}M & \longrightarrow & M & \longrightarrow & \mathrm{Hom}_k(\mathrm{Hom}(M, \mathcal{O}_{x_{m+1}}), \mathcal{O}_{x_{m+1}}) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & \theta\alpha_{(m+1)*}\theta^{-1}\alpha_{m+1}^{*!}N & \longrightarrow & N & \longrightarrow & \mathrm{Hom}_k(\mathrm{Hom}(N, \mathcal{O}_{x_{m+1}}), \mathcal{O}_{x_{m+1}}) & \longrightarrow & 0
\end{array}$$

implies that

$$\theta\alpha_{m+1*}\theta^{-1}\alpha_{m+1}^{*!}M \subset \theta\alpha_{m+1*}\theta^{-1}\alpha_{m+1}^{*!}N \tag{11.57}$$

and thus we may feel free to perform any finite sequence of such pseudo-twists. In particular, for any $q^{\mathbb{Z}}$-orbit in C, we can perform a sequence of downwards pseudo-twists in such a way to ensure that only one point in that orbit is in the support of $M|_C \oplus N|_C$.

Let x_{m+1} be such a point. Then $\mathrm{Hom}(N, \mathcal{O}_{q^d x_{m+1}}) = 0$ for $d \neq 0$ and thus an easy induction using the short exact sequence

$$0 \to \mathcal{O}_{e_{m+1}}(d-1) \to \mathcal{O}_{e_{m+1}}(d) \to \mathcal{O}_{q^{d+1}x_{m+1}} \to 0 \tag{11.58}$$

implies that $\mathrm{Hom}(\alpha_{m+1}^* N, \mathcal{O}_{e_{m+1}}(d)) = 0$ for all d. Since $\alpha_{m+1}^{*!} N$ is a quotient of $\alpha_{m+1}^* N$, we then also have $\mathrm{Hom}(\alpha_{m+1}^* N, \mathcal{O}_{e_{m+1}}(d)) = 0$ for all d. Similarly, for $d \neq 0$,

$$\mathrm{Ext}^1(\mathcal{O}_{q^{d+1}x_{m+1}}, M) \cong \mathrm{Ext}^1(M, \mathcal{O}_{q^d x_{m+1}})^* \cong \mathrm{Ext}^1(M|_C, \mathcal{O}_{q^d x_{m+1}})^* \cong \mathrm{Hom}(\mathcal{O}_{q^d x_{m+1}}, M|_C) = 0 \tag{11.59}$$

and thus $\mathrm{Hom}(\mathcal{O}_{e_{m+1}}(d), \alpha_{m+1}^! M) = 0$ for all d, implying $\mathrm{Hom}(\mathcal{O}_{e_{m+1}}(d), \alpha_{m+1}^{*!} M) = 0$.

Now, consider the morphism $\phi : \alpha_{m+1}^{*!} M \to \alpha_{m+1}^{*!} N$. The direct image is injective with 0-dimensional cokernel, and thus comparing the two spectral sequences for the direct image implies that $\alpha_{(m+1)*} \ker(\phi) = R^1 \alpha_{(m+1)*} \operatorname{coker}(\phi) = 0$, while $R^1 \alpha_{(m+1)*} \ker(\phi) = \alpha_{(m+1)*} \operatorname{coker}(\phi)$ are 0-dimensional. It follows immediately that $c_1(\ker(\phi))$ and $c_1(\operatorname{coker}(\phi))$ are both multiples of e_{m+1}. Since $\alpha_{m+1}^{*!} M$ has no such subsheaf and $\alpha_{m+1}^{*!} N$ has no such quotient, it follows that ϕ is itself injective with 0-dimensional cokernel. Since $\alpha_{m+1}^{*!} N$ is transverse to C with $\chi((\alpha_{m+1}^{*!} N)|_C) < \chi(N|_C)$, the claim follows by induction on $\chi(N|_C)$. □

Remark. It follows from the proof that if $M \subset N$ and the support of $M|_C \oplus N|_C$ hits every $q^{\mathbb{Z}}$-orbit at most once, then $M = N$, as this property is then preserved under taking minimal lifts.

Remark. If C is singular, then x_{m+1} may be fixed by the analogue of $q^{\mathbb{Z}}$; since e_{m+1} is then an anticanonical component, the desired condition on $\alpha_{m+1}^{*!} M$ and $\alpha_{m+1}^{*!} N$ is simply that they both be transverse to C.

Remark. When q is torsion, the claim fails at the base case of the induction: take M to be the kernel of the map from N to a generic 0-dimensional sheaf coming from its support on the center of X.

Let $\mathcal{S}pl_{X_m, X_{m-1}}$ denote the subspace of $\mathcal{S}pl_{X_m}$ classifying simple sheaves such that the derived direct image is a simple sheaf. This is an open subspace (by the same proof as [45, Lem. 7.2]; note that $R^1 \alpha_{m*} M$ is a sheaf on C, so it makes sense to consider its support), and there is a natural morphism $\alpha_{m*} : \mathcal{S}pl_{X_m, X_{m-1}} \to \mathcal{S}pl_{X_{m-1}}$, which certainly surjects on the subspace of pointless sheaves, as it is split by $\alpha_m^{*!}$. (That α_{m*} preserves flatness in α_{m*}-acyclic families follows as in [45, Lem. 7.1])

Theorem 11.23. *[45, Thm. 7.4] The morphism $\alpha_{m*} : \mathcal{S}pl_{X_m, X_{m-1}} \to \mathcal{S}pl_{X_{m-1}}$ is Poisson.*

Proof. The only thing missing to make the proof work is an analogue of [45, Lem. 7.3]. The argument there does not carry over, but we can fix this by realizing that the claim made there has an analogue in the derived category, namely that there is a commutative diagram

$$\begin{array}{ccc} \theta L\alpha_m^* R\alpha_{m*} & \xrightarrow{\theta L\alpha_m^* T} & \theta L\alpha_m^* \theta^{-1} R\alpha_{m*} \\ \downarrow & & \uparrow \\ \theta & \xrightarrow{T} & \mathrm{id} \end{array} \tag{11.60}$$

of triangulated functors, where the right map is

$$\mathrm{id} \to L\alpha_m^! R\alpha_{m*} \cong \theta L\alpha_m^* \theta^{-1} R\alpha_{m*}. \tag{11.61}$$

It suffices to verify this on generators of the derived category, so in particular on the sheaves $\mathcal{O}_{e_m}(-1)$ and $\mathcal{O}_X(D)$ with $D = X_{m-1}(D)$. The first case is trivial, while the second simply states that the pullback of T factors through T.

In particular, we do not need $L\alpha_m^*\alpha_{m*}M$ or $L\alpha_m^!\alpha_{m*}M$ to be sheaves when applying the Lemma, making the proof work even in the case that $\alpha_{m*}M$ is not pointless. □

Corollary 11.24. *The functors α_m^* and $\alpha_m^!$ induce Poisson morphisms $\mathcal{S}pl_{X_{m-1}} \to \mathcal{S}pl_{X_m,X_{m-1}}$.*

Let $\mathcal{S}pl_{X_{m-1},x_m,k}$ denote the locally closed subspace of $\mathcal{S}pl_{X_{m-1}}$ consisting of pointless sheaves such that $\dim \mathrm{Ext}^1(\mathcal{O}_{qx_m}, M) = k$.

Theorem 11.25. *[45] The algebraic space $\mathcal{S}pl_{X_{m-1},x_m,k}$ is a Poisson subspace of $\mathcal{S}pl_{X_{m-1}}$, and $\alpha_m^{*!}$ induces a Poisson morphism $\mathcal{S}pl_{X_{m-1},x_m,k} \to \mathcal{S}pl_{X_m}$.*

Proof. That this is a Poisson subspace follows as in [45, Lem. 7.6] (noting that for simple sheaves, "pointless" is equivalent to "not 0-dimensional", so is both open and closed), and $\alpha_m^{*!}$ preserves flatness for families in $\mathcal{S}pl_{X_{m-1},x_m,k}$, so induces a morphism. That morphism is Poisson since its inverse, α_{m*}, is Poisson. □

If M is a pure 1-dimensional sheaf, then we may define $M^D := R^1 \operatorname{ad} M$, which will again be a pure 1-dimensional sheaf. We then find the following, just as in the commutative case.

Proposition 11.26. *[45, Prop. 7.11] Let M be a pure 1-dimensional sheaf on X_{m-1}. Then $\alpha_m^* M$, $\alpha_m^! M$, $\alpha_m^{*!} M$ are pure 1-dimensional sheaves, and*

$$(\alpha_m^* M)^D \cong \alpha_m^!(M^D)$$
$$(\alpha_m^{*!} M)^D \cong \alpha_m^{*!}(M^D)$$
$$(\alpha_m^! M)^D \cong \alpha_m^*(M^D).$$

Proof. We first note that if N is d-dimensional, then $R^i \operatorname{ad} N = 0$ for $i + d < 2$, by considering $R\operatorname{Hom}(\mathcal{O}_X(-D), R^i \operatorname{ad} N)$ for sufficiently ample D and applying Serre duality. Moreover, we find, since $R\operatorname{ad} R\operatorname{ad} \cong \mathrm{id}$ and $R\operatorname{ad} M$ is supported in degrees 1 and 2 for M 1-dimensional, that $R^1 \operatorname{ad} M$ is at most 1-dimensional and $R^2 \operatorname{ad} M$ is at most 0-dimensional. Moreover, the relevant spectral sequence indicates that $R^2 \operatorname{ad} R^2 \operatorname{ad} M$ is a subsheaf of M, and thus $R^2 \operatorname{ad} M = 0$. It then follows that for a 1-dimensional sheaf M, $R^2 \operatorname{ad} M = 0$ iff M is pure, at which point we conclude that M^D is pure 1-dimensional iff M is pure 1-dimensional. That $\alpha_m^! M$ is pure 1-dimensional follows from the fact that (by adjunction) the direct image of a 0-dimensional subsheaf of $\alpha_m^! M$ is a 0-dimensional subsheaf of M. Since θ preserves dimension, we similarly find that $\alpha_m^* M$ is pure 1-dimensional, and $\alpha_m^{*!} M$ is pure 1-dimensional as a subsheaf of $\alpha_m^! M$.

It remains only to understand how the functors interact with duality. That duality exchanges α^* and $\alpha^!$ follows from the corresponding statement for $R\operatorname{ad}$, for which it suffices to consider line bundles. That the dual of the minimal lift is the minimal lift of the dual then follows from the fact that duality is exact on pure 1-dimensional sheaves. □

The following has the same proof as the commutative case.

Proposition 11.27. *Let $\mathcal{S}pl^1_{X_m}$ denote the subspace of $\mathcal{S}pl_{X_m}$ parametrizing pure 1-dimensional sheaves. Then $-^D$ is an anti-Poisson involution on $\mathcal{S}pl^1_{X_m}$.*

One last thing to consider from [45] is the discussion of rigidity. The following extends directly. Define a relative Euler characteristic $\chi_C(M,N)$ by

$$\chi_C(M,N) = \dim(\mathrm{Hom}(M,N)) - \dim(\mathrm{im}(\mathrm{Ext}^1(M,\theta N) \to \mathrm{Ext}^1(M,N))) + \dim(\mathrm{Hom}(N,M)); \tag{11.62}$$

note that if M is a simple sheaf, then the dimension of the symplectic leaf containing M is $2 - \chi_C(M,M)$.

Proposition 11.28. *One has*

$$\chi_C(M,N) = \chi_C(N,M) = \chi(M,N) + \dim(\mathrm{Ext}^1_C(M|^{\mathbf{L}}_C, N|^{\mathbf{L}}_C)) - \dim(\mathrm{Ext}^2_C(M|^{\mathbf{L}}_C, N|^{\mathbf{L}}_C)). \tag{11.63}$$

Moreover, if M and N are pointless sheaves on X_{m-1}, then $\chi_C(\alpha_m^{*!}M, \alpha_m^{*!}N) = \chi_C(M,N)$.

If M is a simple sheaf with $\chi_C(M,M) = 2$, then M is infinitesimally rigid, in that the moduli space of simple sheaves with the same numerical invariants and derived restriction to C as M is a collection of isolated points (since the tangent space is 0-dimensional everywhere). There are difficulties in extending the classification of rigid 1-dimensional sheaves from the commutative case (the arguments depend too strongly on the notion of "support"). For the explicit example, however, we have something significantly stronger.

Proposition 11.29. *Suppose q is not torsion, and let D be a divisor class on X_m which is admissibly equivalent to a simple root. Then there is at most one isomorphism class of pure 1-dimensional sheaves with first Chern class D, and any such sheaf is simple.*

Proof. The admissible equivalence to a simple root allows us to restrict to the case that D is a simple root. Since there is then a universally ample divisor D' such that $D' \cdot D = 1$, it follows that there is no decomposition $D = D_1 + D_2$ with D_1, D_2 nonzero effective divisor classes. If M, N are two pure 1-dimensional sheaves of Chern class D, then $\chi(M,N) = -2$, and thus at least one of $\mathrm{Hom}(M,N)$ or $\mathrm{Hom}(N,M)$ is nonzero. The lack of a decomposition into effective divisors implies that such a morphism must be injective. Thus the claim will follow if we can show that $\chi(M) = \chi(N)$.

Consider the derived restriction $M|^{\mathbf{L}}_C$. Since M is 1-dimensional, the derived restriction has rank 0; moreover, for any simple root, $D \cdot C_m = 0$, and thus the derived restriction has degree 0. If $M|^{\mathbf{L}}_C$ were not a sheaf, then θM would have a nontrivial subsheaf supported on C, and thus the same would apply to M itself. But this cannot happen; a 0-dimensional subsheaf contradicts purity, while a 1-dimensional subsheaf contradicts indecomposability. Thus $M|^{\mathbf{L}}_C$ is a sheaf of rank and degree 0 on C, implying $M|^{\mathbf{L}}_C = 0$. Since

$$c_1(M|^{\mathbf{L}}_C) = q^{-\chi(M)}\rho(c_1(D)), \tag{11.64}$$

we conclude that there is at most one value of $\chi(M)$ making this work. □

Remark. Note that if q is torsion, we still have rigidity, in that the sheaf is uniquely determined by its numerical invariants; in the torsion case, the Euler characteristic is only fixed modulo the order of q.

We record the following consequence of the proof, which will be useful as a substitute for boundedness results.

Corollary 11.30. *Let q be non-torsion, and let D be a divisor class on X_m. Then there is at most one integer χ such that there exists a pure 1-dimensional sheaf M with $c_1(M) = D$, $\chi(M) = \chi$ and $M|^{\mathbf{L}}_C = 0$.*

Proposition 11.31. *Let D be a (real) root of $W(E_{m+1})$ such that $\rho(D) = q^{-l}$, and suppose that there is a unique isomorphism class of simple sheaves on $X_{\rho;q;C}$ with numerical invariants $(0, D, l)$. Then D is admissibly equivalent to a simple root.*

Proof. Since D is certainly not a simple root, there is a simple root α_i with $\alpha_i \cdot D < 0$. If α_i is not effective, then we may perform the corresponding reflection and proceed by induction (on the height of D, say). Since by assumption D is not admissibly equivalent to a simple root, this process must terminate with α_i effective and different from D.

Let M be a sheaf with the given invariants (if no such sheaf exists, there is nothing to prove), and let M_0 be a pure 1-dimensional sheaf with Chern class α_i. Then $\chi(M_0, M) = -\alpha_i \cdot D > 0$, so that at least one of $\mathrm{Hom}(M_0, M)$ or $\mathrm{Hom}(M, M_0)$ must be nonzero. If q has order $r < \infty$, we may choose a divisor class D' such that $D' \cdot D = 0$, $D' \cdot \alpha_i \neq 0$, and note that twisting by nrD' for $n \gg 0$ or $n \ll 0$ as appropriate will give a nonisomorphic sheaf (by preventing a map from or to M_0). We may thus assume q non-torsion at this point. Then M_0 has no nonzero proper subsheaves, and thus any composition of nonzero maps $M \to M_0 \to M$ would be a nontrivial endomorphism of M, so that at most one of the two spaces can be nonzero.

If $\mathrm{Hom}(M_0, M) \neq 0$ (the case $\mathrm{Hom}(M, M_0) \neq 0$ is analogous), then let N be the complex fitting into the distinguished triangle

$$R\,\mathrm{Hom}(M_0, M) \otimes M_0 \to M \to N \to \quad ; \tag{11.65}$$

that is N is the image of M under the appropriate spherical twist functor [57] (which we call a *reflection functor* to avoid confusion with twisting by line bundles). The map $\mathrm{Hom}(M_0, M) \otimes M_0 \to M$ is injective, and thus N is a sheaf, and $\mathrm{Hom}(M_0, N) = 0$. Since M_0 is a spherical object, the above process is reversible, so in particular $\mathrm{Hom}(N, N) \cong \mathrm{Hom}(M, M)$ and N is simple. If $\mathrm{Hom}(N, M_0) = 0$, then repeating the construction gives a simple sheaf M' with the same Chern class as M but satisfying $\mathrm{Hom}(M_0, M') = 0$, so that $M \ncong M'$. On the other hand, if $\mathrm{Hom}(N, M_0) \neq 0$, then $c_1(N)$ is a root of lower height than D, and thus by induction there exists a nonisomorphic simple sheaf N' with the same numerical invariants. If $\mathrm{Hom}(N', M_0) = 0$, then repeating the above construction gives a sheaf $M' \ncong M$; if $\mathrm{Hom}(M_0, N') = 0$, then we instead apply the inverse:

$$M' \to N' \to R\,\mathrm{Hom}_k(R\,\mathrm{Hom}(N', M_0), M_0) \to \quad . \tag{11.66}$$

□

Remark. It follows (for q non-torsion) that given any such sheaf M, there is a sequence of reflection functors taking M to a non-isomorphic sheaf M' with the same invariants. It also follows from the above induction that any simple 1-dimensional sheaf with Chern class a root is related via a sequence of reflection functors (preserving sheaf-hood) and admissible reflections to a sheaf with Chern class a simple root.

11.2 Hilbert schemes

As we mentioned above, one of the difficulties in dealing with moduli spaces on these noncommutative surfaces is the absence of good boundedness results. As a result, we can only give partial results even on the existence of (quasi)projective moduli spaces of (semi)stable sheaves. Luckily, the cases we *can* deal with include some particularly nice examples. In the present section, we consider the case of Hilbert schemes of points; in the next, we will discuss moduli spaces of 1-dimensional sheaves.

We should first note the following general statements. By D-stable or D-semistable for an ample divisor class D, we mean the obvious analogue of the usual Gieseker notions from the commutative setting, in terms of the Hilbert polynomial $p_M(t) := \chi(\mathcal{O}_X(-tD), M)$.

Proposition 11.32. *For any ample divisor class D, the moduli functor of D-stable sheaves on $X_{\rho;q;C}$ is separated.*

Proof. The standard proof in the commutative case carries over directly: given two families of stable sheaves on a discrete valuation ring with isomorphic generic fibers, semicontinuity gives a nonzero morphism between the special fibers, and stability forces that to be an isomorphism. □

Proposition 11.33. *For any ample divisor class D, the moduli functor of D-semistable sheaves on $X_{\rho;q;C}$ is proper.*

Proof. Again, the commutative argument carries over. Given a semistable sheaf over the open point of a discrete valuation ring, we may extend in any way to a family over the ring itself, and then proceed as in [31, 34, 26] to modify this to have semistable special fiber. (Here, we need the fact that if M is a d-dimensional sheaf, then there is a natural injection $M \mapsto R^{2-d}\operatorname{ad} R^{2-d}\operatorname{ad} M$, but this follows immediately from the isomorphism $M \cong R\operatorname{ad} R\operatorname{ad} M$ using the appropriate spectral sequence.)

It remains to show that the moduli functor is separated; in other words, given two semistable families M and N over a discrete valuation ring with isomorphic generic fibers, we need to show that the special fibers are S-equivalent. The isomorphism of generic fibers implies the existence of an injective morphism $\pi^k N \to M$ for some k; replacing N by the image of $\pi^k N$ reduces to the case $\pi^l M \subset N \subset M$.

Thus suppose M and N are semistable sheaves with $\pi^l M \subset N \subset M$, and let $M' := N + \pi M$. Since $\pi^{l-1} M' \subset N \subset M'$, it will suffice to show that M' is semistable with special fiber S-equivalent to that of M. If $M' = M$ or $M' = \pi M$, then the special fibers are isomorphic. Otherwise, we have a four-term exact sequence

$$0 \to \pi M/\pi M' \to M'/\pi M' \to M/\pi M \to M/M' \to 0 \tag{11.67}$$

in which the middle morphism is neither 0 nor an isomorphism. The image of the middle morphism is simultaneously a subsheaf of $M/\pi M$ and a quotient of $N/\pi N$, and thus has the same reduced Hilbert polynomial; it follows that the same applies to $M/M' \cong \pi M/\pi M'$. But this immediately implies that $M'/\pi M'$ and $M/\pi M$ are S-equivalent, as required. □

In the commutative setting, the Hilbert scheme of points is most naturally dealt with as a Quot scheme. This behaves badly in the noncommutative setting; instead, one simply (following [36]) considers sheaves with the same numerical invariants as an ideal sheaf of n points.

Definition 11.2. The *Hilbert scheme of n points on $X_{\rho;q;C}$*, denoted $\operatorname{Hilb}^n(X_{\rho;q;C})$ is the moduli space of pure 2-dimensional sheaves M with numerical invariants $(1, 0, 1-n)$.

Remark. The notation reflects the fact that when $q = 1$, this is just the Hilbert scheme of n points on $X_{\rho;1;C}$. Indeed, any rank 1 sheaf M on a commutative surface is contained in a line bundle with the same Chern class, namely the double dual M^{**}.

Proposition 11.34. *Suppose M is a pure 2-dimensional sheaf of rank 1. Then M is simple.*

Proof. Suppose M is not simple, and let $\phi : M \to M$ be a nonscalar endomorphism of M. Since $\mathrm{rank}(M) > 0$, we find that $M|_C \neq 0$, and thus ϕ induces an endomorphism of $M|_C$. In particular, some linear combination of ϕ and $\mathrm{id}(M)$ will fail to be surjective on $M|_C$, and thus fail to be surjective. It follows (by Euler characteristic considerations) that the kernel of ϕ is also nontrivial. Purity of M implies that $\mathrm{rank}(\ker(\phi)) > 0$ and $\mathrm{rank}(\mathrm{im}(\phi)) > 0$; since $\mathrm{rank}(\ker(\phi))+\mathrm{rank}(\mathrm{im}(\phi)) = 1$, this is impossible. □

Lemma 11.35. *If $n < 0$, then $\mathrm{Hilb}^n(X_{\rho;q;C})$ is empty, while $\mathrm{Hilb}^0(X_{\rho;q;C})$ consists of the single point $\mathcal{O}_X$.*

Proof. If I is pure 2-dimensional with numerical invariants $(1, 0, c)$ with $c \geq 1$, then I has a global section since Chern class considerations forbid a morphism from I to $\theta\mathcal{O}_X$. Any global section is injective, and the cokernel is a 0-dimensional sheaf, N say. Since any 0-dimensional sheaf satisfies $\mathrm{Ext}^1(N, \mathcal{O}_X) \cong \mathrm{Ext}^1(\mathcal{O}_X, \theta N)^* = 0$, the extension is split, contradicting the pure 2-dimensionality of I unless $N = 0$. □

This gives the following characterization of line bundles on $X_{\rho;q;C}$.

Corollary 11.36. *If M is a pure 2-dimensional sheaf on $X_{\rho;q;C}$ of rank 1, then $\chi(M, M) \leq 1$, with equality iff M is a line bundle.*

Proof. This follows from the Lemma and the fact that $M(-c_1(M))$ is a pure 2-dimensional sheaf with numerical invariants $(1, 0, \frac{1+\chi(M,M)}{2})$. □

Remark. Substituting in the formula for χ gives the bound in the form

$$\chi(M) \leq 1 + \frac{c_1(M) \cdot (c_1(M) + C_m)}{2}. \tag{11.68}$$

This implies some useful boundedness results for globally generated sheaves.

Lemma 11.37. *If M is a 1-dimensional sheaf such that there is a surjective morphism $\mathcal{O}_X \to M$, then $\chi(M) \geq -\frac{c_1(M)\cdot(c_1(M)-C_m)}{2}$.*

Proof. The kernel I of such a morphism is a pure 1-dimensional sheaf with $c_1(I) = -c_1(M)$, $\chi(I) = 1 - \chi(M)$, so that the result follows immediately from Corollary 11.36. □

Remark. By the proof, the bound is tight iff M is the quotient of a morphism $\mathcal{O}_X(-D) \to \mathcal{O}_X$. In the commutative case, the same quantity is an upper bound if we assume that M is pure, but for q non-torsion, there is no upper bound in general. Indeed, on X_0 with q non-torsion, any sheaf of the form $\mathcal{O}_{s-f}(d)$, $d \geq 0$, is generated by any of its global sections, but has Euler characteristic exceeding the lower bound by d.

Proposition 11.38. *For every divisor class D, there is a bound $B(D)$ such that if M is a globally generated sheaf on any $X_{\rho;q;C}$ with numerical invariants (r, D, l), then $l \geq r + B(D)$. Moreover, if D is not effective, then there are no such sheaves.*

Proof. If $D = 0$, take $B(D) = 0$. Otherwise, let M be such a sheaf. If M has a 0-dimensional subsheaf T, we can lower $\chi(M)$ without changing $c_1(M)$ by taking the quotient of M by T. Similarly, if $\mathrm{rank}(M) > 0$, then there is an injective map $\mathcal{O}_X \to M$, and we can take the quotient without changing $\chi(M) - \mathrm{rank}(M)$. Assume, therefore, that M is pure 1-dimensional (implying

$c_1(M)$ effective!), and let $\phi : \mathcal{O}_X \to M$ be a global section of M. Then $D' = c_1(\operatorname{im}\phi) \neq 0$, so $c_1(\operatorname{coker}\phi) < c_1(M)$, and thus by induction we have

$$\chi(M) \geq -\frac{D' \cdot (D' - C_m)}{2} + B(D - D'). \tag{11.69}$$

By the proof of Lemma 10.21, there are only finitely many possible values for D' even as we allow the surface to vary, and thus we may take the minimum

$$B(D) := \min_{D'} \left(-\frac{D' \cdot (D' - C_m)}{2} + B(D - D') \right), \tag{11.70}$$

where D' ranges over nonzero divisors such that both D' and $D - D'$ are in the dual to the universal nef cone. □

Remark. If we expand the above description to the full decomposition of D, we can simplify the resulting expression to obtain

$$B(D) = -\frac{D \cdot (D - C_m)}{2} + \min_{D_1 + \cdots + D_l = D} \sum_{1 \leq i < j \leq l} D_i \cdot D_j, \tag{11.71}$$

where the minimum is over D_i in the dual to the universal nef cone. Note that if $D_i \cdot D_j \geq 0$ for some $i < j$, then we can merge those two components without increasing the value, and thus only those decompositions for which $D_i \cdot D_j < 0$ for all $i < j$ need be considered.

Corollary 11.39. *Let $\mathcal{M}$ be a set of globally generated sheaves on surfaces X_m such that the corresponding set of numerical invariants is finite. Then the set of numerical invariants of globally generated subsheaves of sheaves in $\mathcal{M}$ is also finite.*

Proof. Let $M \in \mathcal{M}$ and let N be a subsheaf. Since $0 \leq \operatorname{rank}(N) \leq \operatorname{rank}(M)$, there are only finitely many possible values for $\operatorname{rank}(N)$. Moreover, since both N and M/N are globally generated, we find that $c_1(N)$ and $c_1(M) - c_1(N)$ are both effective. Then Lemma 10.21 tells us that there are only finitely many possible values for $c_1(N)$. We then have the bounds

$$\operatorname{rank}(N) + B(c_1(N)) \leq \chi(N), \quad \text{and} \quad \operatorname{rank}(M) - \operatorname{rank}(N) + B(\operatorname{rank}(M) - c_1(N)) \leq \chi(M) - \chi(N), \tag{11.72}$$

giving a finite range of Euler characteristics for each choice of $c_1(N)$. □

Lemma 11.40. *If M is a pure 1-dimensional sheaf of Chern class df on X_0 or X_0', then there is a filtration M_i of M such that each M_{i+1}/M_i is pure 1-dimensional of Chern class f and such that*

$$\chi(M_1) \geq \chi(M_2/M_1) \geq \cdots \chi(M/M_{d-1}). \tag{11.73}$$

Proof. Since M is pure 1-dimensional, there is some twist of M that has no global sections; the claim is invariant under twisting, so we may assume $h^0(M(-s)) = 0$. Since $\chi(M(ns)) > 0$ for $n \gg 0$, there is some smallest n_1 such that $M(n_1 s)$ has a global section, and again we may assume $n_1 = 0$. Now, let N be the image of the natural map $\mathcal{O}_X \otimes_k \operatorname{Hom}(\mathcal{O}_X, M) \to M$. If N has Chern class $d_1 f$, then we have $\chi(N) \geq B(d_1 f)$. An easy induction shows that $B(df) = d$ for all d, and thus $\chi(N) \geq d_1$. Since $h^0(N(-s)) = 0$, we also have $\chi(N(-s)) \leq 0$, and this implies $\chi(N) = d_1$. It follows that N is a direct sum of quotients $\mathcal{O}_X/\mathcal{O}_X(-d'f)$, and is thus the pullback of a 0-dimensional sheaf on $\mathbb{P}^1$. Taking a filtration of that 0-dimensional sheaf with degree 1 subquotients gives a filtration of N in which each subquotient is the pullback of a point.

In particular, N is acyclic, and thus M/N has no global sections, so is pure 1-dimensional. By induction, M/N has a filtration as desired, and $h^0(M/N) = 0$ implies that the Euler characteristic of the bottom subsheaf is nonnegative, and thus strictly less than the subquotients of the filtration of N. Gluing the two filtrations gives the desired result. □

We can also control acyclicity of twists.

Lemma 11.41. *If M is a globally generated coherent sheaf, then for any nef divisor class D with $D \cdot C_m > 0$ and $c_1(M) - D - C_m$ ineffective, $M(D)$ is acyclic.*

Proof. If $n = 1$, $\operatorname{rank}(M) = 1$, this is immediate from $M \cong \mathcal{O}_X$. If $n = 1$, $\operatorname{rank}(M) = 0$, let I be the kernel of the given global section. Since D is nef, $\operatorname{Ext}^p(\mathcal{O}_X(-D), \mathcal{O}_X) = 0$ for $p > 0$, and thus we find $\operatorname{Ext}^2(\mathcal{O}_X(-D), M) = 0$ and

$$\operatorname{Ext}^1(\mathcal{O}_X(-D), M) \cong \operatorname{Ext}^2(\mathcal{O}_X(-D), I) \cong \operatorname{Hom}(I, \theta\mathcal{O}_X(-D))^*. \tag{11.74}$$

A nonzero morphism $I \to \theta\mathcal{O}_X(-D)$ would be injective (both sheaves have only rank 1 subsheaves) and the cokernel would be a 1-dimensional sheaf of ineffective Chern class $c_1(M) - C_m - D$.

For general n, choose one global section, and let I_1, M_1 be the corresponding kernel and image. By Corollary 11.39, there is only a finite set of possibilities for the numerical invariants of M_1, and thus the argument of the previous paragraph gives a uniform bound for I_1, M_1, while induction gives a uniform bound for I/I_1, M/M_1. □

We will need some boundedness results on the Hilbert scheme.

Lemma 11.42. *For $m > 0$, if I is a pure 2-dimensional sheaf on X_m with numerical invariants $(1, 0, 1 - n)$, then $\alpha_{m*}I(-ne_m)$ is a pure 2-dimensional sheaf on X_{m-1} with numerical invariants $(1, 0, 1 - n(n + 3)/2)$, and $I(-ne_m) \cong \alpha_m^{*!}\alpha_{m*}I(-ne_m)$.*

Proof. Since I is pure 2-dimensional, $\operatorname{Hom}(\mathcal{O}_{e_m}(-1), I(-ne_m)) = 0$. On the other hand, if there were a morphism $I(-ne_m) \to \mathcal{O}_{e_m}(-1)$, then the cokernel would be 0-dimensional, and thus the kernel would be a pure 2-dimensional sheaf with numerical invariants $(1, -(n+1)e_m, d-n(n+3)/2)$ with $d > 0$, violating Corollary 11.36. The claim follows from Lemma 11.17. □

Lemma 11.43. *If I is a pure 2-dimensional sheaf on X_0 or X_0' with numerical invariants $(1, 0, 1 - n)$, then $I(nf + D)$ is acyclic and globally generated for any nef divisor class D such that $D \cdot f \geq n$.*

Proof. Since $\chi(I(nf)) = 1$ and $\operatorname{Ext}^2(\mathcal{O}_X(-nf), I) \cong \operatorname{Hom}(I, \theta\mathcal{O}_X(-nf)) = 0$ by Chern class considerations, there is a morphism $\mathcal{O}_X(-nf) \to I$, giving a short exact sequence

$$0 \to \mathcal{O}_X(-nf) \to I \to M \to 0. \tag{11.75}$$

Since $M(nf)$ has no global sections, it follows from Lemma 11.40 that it has a filtration in which each subquotient has numerical invariants $(0, f, l)$ for $l \leq 0$. It follows that each $l \leq n$, and thus $M(ds + d'f)$ is acyclic and globally generated for $d \geq n$. Since $\mathcal{O}_X(-nf + ds + d'f)$ is acyclic for $0 \leq d \leq d' - n$, we conclude that $I(ds + d'f)$ is acyclic and globally □

This allows us to fill in the missing boundedness result for Theorem 10.37.

Lemma 11.44. *For any choice of numerical invariants (r, D, l) and any divisor class D_a, there is an integer B such that for any X_m on which D_a is ample and any short exact sequence*

$$0 \to I \to \mathcal{O}_X^n \to M \to 0 \tag{11.76}$$

for which M has numerical invariants (r, D, l), the sheaves $I(bD_a)$ and $M(bD_a)$ are acyclic and globally generated.

Proof. The same argument as in the proof of Lemma 11.41 allows us to reduce to the case $n = 1$, $r = 0$. Then I has numerical invariants $(1, -D, 1 - l)$, and thus $I(D)$ has invariants $(1, 0, 1 - l - D \cdot (D - C_m)/2)$. For $m > -1$, we may inductively apply Lemmas 11.42 and 11.43 to construct a (universal) divisor class $D_{m,n}$ such that $I(D + D_{m,n})$ is acyclic and globally generated. Then Lemma 11.41 extends this to a translate of the relevant nef cone, and thus to multiples of D_a.

For $m = -1$, we simply apply Proposition 10.36 to the Hilbert scheme, which is projective by [36, Thm. 8.11]. □

Remark. Note that since D is effective, the proof of Lemma 10.21 tells us that there is only a finite set of possible numerical invariants for any given Hilbert polynomial.

We can now prove that the Hilbert scheme is projective.

Lemma 11.45. *The Hilbert scheme* $\mathrm{Hilb}^n(X_{\rho;q;C})$ *is a smooth, proper, Poisson algebraic space of dimension* $2n$.

Proof. That it is a Poisson algebraic space is inherited from $\mathcal{S}pl_{X_{\rho;q;C}}$, and smoothness follows from the fact that $\mathrm{Ext}^2(M, M) \cong \mathrm{Hom}(M, \theta M)^* = 0$, so deformations are unobstructed, while properness follows from the fact that a sheaf with the given numerical invariants is semistable iff it is stable iff it is pure 2-dimensional. The dimension follows from $\chi(M, M) = 1 - 2n$ and the fact that M is simple. □

Theorem 11.46. *The Hilbert scheme* $\mathrm{Hilb}^n(X_{\rho;q;C})$ *is a smooth, irreducible projective scheme.*

Proof. For X_{-1}, this was shown in [36, Thm. 8.11]. Otherwise, it follows from Lemmas 11.42 and 11.43 that there is a divisor class D such that for any I (on any X_m) with numerical invariants $(1, 0, 1 - n)$, $I(D)$ is acyclic and globally generated. Since Corollary 11.39 implies that there are only finitely many possible values for the numerical invariants of a globally generated subsheaf of $I(D)$, it follows from [26, Lems. 4.4.5,4.4.6] that there is a linearization of the action of $\mathrm{GL}(\mathrm{Hom}(\mathcal{O}_X(-D), I))$ on the Quot scheme such that every point corresponding to I is GIT stable. (Note that since I has rank 1, any subsheaf still satisfies the relevant bound, even if it has more global sections than expected.) But then $\mathrm{Hilb}^n(X_{\rho;q;C})$ is a subscheme of the GIT quotient, and is therefore quasiprojective, so projective.

It remains to show irreducibility. Here, we use the fact that since the moduli problem is unobstructed, the moduli space remains smooth even as we vary $X_{\rho;q;C}$ in a family. (Moreover, our choices of ample divisor class could be made in a universal fashion, so the construction continues to work as we allow the surface to vary.) In particular, we may consider the family with base $\mathrm{Pic}^0(C)$ obtained by keeping ρ and C fixed but allowing q to vary. The argument of [36, Prop. 8.6] then applies to say that if any fiber is connected, then all fibers are connected. Since for $q = 1$ this is just the usual Hilbert scheme, connectedness (and thus irreducibility since $\mathrm{Hilb}^n(X_{\rho;q;C})$ is smooth) follows immediately. □

Remark 1. The quasiprojective subscheme on which $I|_C \cong \mathcal{O}_C$ is symplectic, and can be thought of as an analogue of Calogero-Moser space, see the introduction to [36].

In the case $n = 1$, $\mathrm{Hilb}^1(X_{\rho;q;C})$ is a surface, so we ought to identify precisely which projective surface results.

Proposition 11.47. *For any* ρ, q, C, $\mathrm{Hilb}^1(X_{\rho;q;C}) \cong X_{\rho;1;C}$.

Proof. The key idea is that if M is a pure 2-dimensional sheaf with the given numerical invariants, then $R\operatorname{Hom}(\mathcal{O}_X, M) = 0$. Indeed, $\operatorname{Hom}(\mathcal{O}_X, M) = 0$ and $\operatorname{Hom}(M, \theta\mathcal{O}_X) = 0$ since such a morphism would need to be injective, and the cokernel would have ineffective Chern class. Since $\chi(\mathcal{O}_X, M) = \chi(M) = 0$, this implies that $\operatorname{Ext}^1(\mathcal{O}_X, M) = 0$ as well, giving $R\operatorname{Hom}(\mathcal{O}_X, M) = 0$.

In other words, M is an object in the subcategory $\mathcal{N}_{\rho;q;C} \subset D^b \operatorname{coh} X_{\rho;q;C}$. We may thus use the functor κ_q to move this to $\mathcal{N}_{\rho;1;C}$. We claim that $\kappa_q M$ is a pure 2-dimensional sheaf, and conversely that for any point $x \in X_{\rho;1;C}$, $\kappa_q^{-1} I_x$ is a pure 2-dimensional sheaf. Moreover, if $x \in C$, then $\kappa_q^{-1} I_x$ is the ideal sheaf of qx.

We first consider the case $m = -1$. In this case, we similarly find that $\operatorname{Hom}(\mathcal{O}_X(\pm 1), M) = \operatorname{Ext}^2(\mathcal{O}_X(\pm 1), M) = 0$. This implies that M has a (minimal) resolution of the form

$$0 \to \mathcal{O}_X(-2) \to \mathcal{O}_X(-1)^2 \to M \to 0; \tag{11.77}$$

conversely, any pair of linearly independent morphisms in $\operatorname{Hom}(\mathcal{O}_X(-2), \mathcal{O}_X(-1))$ will have pure 2-dimensional quotient. Since this condition is preserved by κ_q, the claim follows. (Note also that the moduli space in this case is $\operatorname{Gr}(2, \operatorname{Hom}(\mathcal{O}_X(-2), \mathcal{O}_X(-1)))) \cong \mathbb{P}^2$.) The claim in the case of an ideal sheaf of a point of C follows from the fact that κ_q respects restriction to C, and the torsion subsheaf of the restriction to C of the ideal sheaf of a point $x \in C$ is the structure sheaf of x/q. A similar argument applies to $X_{\eta,\eta';q;C}$.

For the remaining cases, let $\mathcal{O}_e(-1)$ denote either $\mathcal{O}_{e_m}(-1)$ or $\mathcal{O}_s(-1)$ as appropriate, with α_*, α^* the corresponding functors associated to the blowdown. If $\operatorname{Hom}(M, \mathcal{O}_e(-1)) = 0$, then M is α_*-acyclic, and $\alpha_* M$ is pure 2-dimensional, with $M \cong \alpha^* \alpha_* M$, and thus the claim follows by induction.

If $\operatorname{Hom}(M, \mathcal{O}_e(-1)) \neq 0$, then the image is $\mathcal{O}_e(-d)$ for some $d \geq 1$, but then applying Corollary 11.36 as in the proof of Lemma 11.42 tells us that $d = 1$ and that we have a short exact sequence

$$0 \to \mathcal{O}_X(-e) \to M \to \mathcal{O}_e(-1) \to 0. \tag{11.78}$$

Since $\mathcal{O}_X(-e) \cong \alpha^* \alpha_* \mathcal{O}_X(-e)$ and $\alpha_* \mathcal{O}_X(-e)$ is the ideal sheaf of a point, this structure is preserved by κ_q and κ_q^{-1} as required. □

In general, of course, we expect $\operatorname{Hilb}^n(X_{\rho;q;C})$ to be a nontrivial deformation of $\operatorname{Hilb}^n(X_{\rho;1;C})$. It is, however, closely related to commutative Hilbert schemes; this will, in particular, allow us to prove that $\operatorname{Hilb}^n(X_{\rho;q;C})$ is rational.

The basic idea in the construction is that if D is a divisor class such that $\chi(I(D)) = 0$, then we can expect an open subset on which $I(D) \in \mathcal{N}_{\rho;q;C}$, allowing us to apply κ_q and then untwist. Of course, this will in general only give us an object in the derived category, but we can at least hope that this will generically be a sheaf. On one side, this is well-behaved.

Lemma 11.48. *Let S be a locally noetherian scheme, and let $\phi : M \to N$ be a morphism of S-flat families of sheaves on $X_{\rho;q;C}$. Then the set of points of S for which the corresponding fiber of ϕ is surjective is open.*

Proof. Suppose first that $m > 0$ and (as we may) that S is noetherian. By the proof of Proposition 10.36, there is an integer D such that for $d \geq D$, $\theta^{-d} M$ and $\theta^{-d} N$ are both acyclic for α_{m*}. By induction, for each $d \geq D$, there is an open subset U_d on which $\alpha_{m*} \theta^{-d} \phi$ is surjective. Now, if $\alpha_{m*} \theta^{-d} \phi$ is surjective, then the standard spectral sequence gives $R\alpha_{m*} \operatorname{coker}(\theta^{-d} \phi) = 0$, and thus $\operatorname{coker}(\theta^{-d} \phi) \cong \mathcal{O}_{e_m}(-1)^a$ for some a. But then $\operatorname{coker}(\theta^{-d-1} \phi) \cong \mathcal{O}_{e_m}^a$ and $\alpha_{m*} \operatorname{coker}(\theta^{-d-1} \phi) \neq 0$. In other words, ϕ is surjective at every point in $U_d \cap U_{d+1}$. But any surjective fiber is in U_d for sufficiently large d, and thus the desired open subset can be written as the union $\cup_{d \geq D} (U_d \cap U_{d+1})$.

For $m \leq 0$, let D_a be an ample divisor as in Corollary 10.33 or 10.34 as appropriate. Let U_l be the open subset of S where N is $(-l, D_a)$-regular and the induced map

$$\mathrm{Hom}(\mathcal{O}_X(-lD_a), M) \to \mathrm{Hom}(\mathcal{O}_X(-lD_a), N) \tag{11.79}$$

is surjective. Then ϕ is surjective on the fiber $v \in S$ iff $v \in U_l$ for some l, giving the desired result. □

Lemma 11.49. *Let S be a locally noetherian scheme, and let $M^{\cdot}$ be an S-flat family of bounded coherent complexes on $X_{\rho;q;C}$. Then the set of points of S for which the corresponding fiber of $M^{\cdot}$ only has cohomology in degree 0 is an countable intersection of open subsets.*

Proof. By the previous lemma, we know that the set where $M^{\cdot}$ has nonpositive cohomology is open. On the other hand, the set where $M^{\cdot}$ has only nonnegative cohomology can be expressed as the countable intersection of open subsets $\mathrm{Ext}^i(\mathcal{O}_X(D), M^{\cdot}) = 0$ for $i < 0$, $D \in \mathrm{Pic}(X)$. □

Remark. In particular, if S is integral, then the generic fiber is a sheaf as long as at least one fiber is sheaf.

Lemma 11.50. *For any n, there is a birational map $\mathrm{Hilb}^n(X_{\rho;q;C}(-(n-1)f)) \dashrightarrow \mathrm{Hilb}^n(X_{\rho;1;C})$ given by*

$$I \mapsto \kappa_q(I((n-1)f))(-(n-1)f) \tag{11.80}$$

on the nonempty intersection of open sets where $I((n-1)f) \in \mathcal{N}_{\rho;q;C}$ and $\kappa_q(I((n-1)f))$ is a pure 2-dimensional sheaf.

Proof. Since all of the conditions we are imposing are open (or intersections of open conditions), we need only show that they hold generically, or equivalently that there is *some* point $I \in \mathrm{Hilb}^n(X_{\rho;q;C}(-(n-1)f))$ for which everything holds. Let $y_1, \ldots, y_n$ be sufficiently general points of C, and consider the corresponding blowup X_{m+n}. The direct image on $X_m(-(n-1)f)$ of the sheaf $\mathcal{O}_X(-e_{m+1} - \cdots - e_{m+n})$ on $X_{m+n}(-(n-1)f)$ corresponds to a point of $\mathrm{Hilb}^n(X_{\rho;q;C}(-(n-1)f))$. Since $R\Gamma(\mathcal{O}_X((n-1)f - e_{m+1} - \cdots - e_{m+n})) = 0$ and κ_q respects blowups and the action of $W(E_{m+n+1})$, we reduce to showing that

$$\kappa_q(\mathcal{O}_X((n-1)f - e_1 - \cdots - e_n)) \cong \mathcal{O}_X((n-1)f - e_1 - \cdots - e_n) \tag{11.81}$$

on the generic n-point blowup of X_0. In fact, using reflections in $W(D_{m+n})$ along with an elementary transformation if n is odd, we find that this is equivalent to showing

$$\kappa_q(\mathcal{O}_X(-f + e_1 + \cdots + e_n)) \cong \mathcal{O}_X(-f + e_1 + \cdots + e_n), \tag{11.82}$$

which follows from the $n = 0$ case and the compatibility of κ_q with $\alpha_m^!$. □

Since Hilbert schemes of commutative rational surfaces are rational, we conclude the following.

Corollary 11.51. *The Hilbert scheme $\mathrm{Hilb}^n(X_{\rho;q;C})$ is rational.*

It will be helpful below to consider a generalization of the above construction.

Proposition 11.52. *Let D be a nef divisor on $X_{\rho;q;C}$ with $D \cdot (D + C_m)/2 = n - 1$, $D \cdot C_m > 0$. Then the generic point $I \in \mathrm{Hilb}^n(X_{\rho;q;C}(-D))$ satisfies $I(D) \in \mathcal{N}_{\rho;q;C}$ and $(\kappa_q I(D))(-D) \in \mathrm{Hilb}^n(X_{\rho;1;C})$.*

Proof. It will be convenient to prove instead that the generic point $I \in \mathrm{Hilb}^n(X_{\rho;1;C})$ is such that $I(D) \in \mathcal{N}_{\rho;q;C}$ and $\kappa_q^{-1} I(D)$ is a pure 2-dimensional sheaf; this will still give us a birational map, and thus the twist of the image of the generic point in $\mathrm{Hilb}^n(X_{\rho;1;C})$ will be the generic point in $\mathrm{Hilb}^n(X_{\rho;q;C}(-D))$ as required.

We may in addition assume that D is universally nef. Moreover, we have already dealt with the case $D = (n-1)f$; we also find that the conclusion holds for the non-nef divisors $D \in \{-s-f, -2s-2f\}$ on X_{-1} and $D = -s + lf$ on either X_0, since in each of those cases $n = 0$ so $I \cong \mathcal{O}_X$. Similarly, if $D \cdot e_m = 0$, then the generic I is an inverse image of some I on X_{m-1}, and thus we may reduce to the blown down surface. The same argument allows us to blow down from X_1 if $D \cdot (f - e_1) = 0$.

Having blown down when possible, we find that $D - C_m$ is either universally nef or one of the above non-nef special cases, and thus by induction the claim holds for $D - C_m$. It suffices to prove that for *some* I, $R\Gamma I(D) = 0$ and $\kappa_q^{-1} I(D)$ is a pure 2-dimensional sheaf, since the conditions are intersections of open conditions. Thus let Z' be a generic $D \cdot (D - C_m)/2 + 1$-point subscheme of $X_{\rho;1;C}$, and let $I_{Z'}$ be its ideal sheaf. By induction, $\theta \kappa_q^{-1} I_{Z'}(D) \cong \kappa_q^{-1} I_{Z'}(D - C_m)$ is a pure 2-dimensional sheaf. Let Z'' be a generic $D \cdot C_m$-point subscheme of C, and consider the subscheme $Z := Z' \cup Z''$. We have a non-split short exact sequence

$$0 \to \theta I_{Z'} \to I_Z \to L \to 0 \tag{11.83}$$

where L is the ideal sheaf of Z'' on C. In particular, $L(D)$ is a degree 0 line bundle, and genericity gives $L(D) \in \mathcal{N}_{\rho;1;C}$, so the same holds for $I_Z(D)$. Twisting by D and applying κ_q^{-1} gives the short exact sequence

$$0 \to \kappa_q^{-1} \theta I_{Z'}(D) \to \kappa_q^{-1} I_Z(D) \to L(D) \otimes q^{-1} \to 0, \tag{11.84}$$

where we used the fact that, by genericity, $L(D) \otimes q^{-1} \in \mathcal{N}_{\rho;1;C}$. It follows that $\kappa_q^{-1} I_Z(D)$ is a sheaf, and it remains only to show that it is pure 2-dimensional.

To see this, note that since $I_{Z'}(D)$ is pure 2-dimensional, the torsion subsheaf of $I_Z(D)$ must be a subsheaf of the line bundle $L(D) \otimes q^{-1}$, so if $I_Z(D)$ is not pure 2-dimensional, its torsion subsheaf is a line bundle L' on C. Then the long exact sequence for restriction to C starts

$$0 \to \theta L' \to \theta L(D) \otimes q^{-1} \to \theta I_{Z'}(D)|_C \otimes q^{-1} \tag{11.85}$$

Since $I_{Z'}(D)|_C$ is torsion-free, we conclude that $L' \cong L(D) \otimes q^{-1}$. But this contradicts the non-split condition! □

We could also make the Euler characteristic 0 by twisting down rather than twisting up. Such cases can be easily understood, at least when q is non-torsion, via the following fact; compare [36, Thm. 8.11(3)].

Proposition 11.53. *If $I \in \mathrm{Hilb}^n(X_{\rho;q;C})$ with $I|_C \cong q^{-n}$ and $n < |\langle q \rangle|$, then $R \operatorname{ad} I$ is a sheaf and $\theta^{-1} \operatorname{ad} I \in \mathrm{Hilb}^n(X_{\rho;1/q;C})$.*

Proof. We recall that $R^2 \operatorname{ad} I$ is ≤ 0-dimensional and $R^1 \operatorname{ad} I$ is ≤ 1-dimensional. In fact, if either bound were tight, then the isomorphism $R \operatorname{ad} R \operatorname{ad} I \cong I$ would produce a < 2-dimensional subsheaf of I. We thus conclude that $R^2 \operatorname{ad} I = 0$ and $R^1 \operatorname{ad} I$ is ≤ 0-dimensional. Similarly, if $\operatorname{ad} I$ had a < 2-dimensional subsheaf, this would prevent $R \operatorname{ad} R \operatorname{ad} I \cong I$ from being a sheaf, and thus $\operatorname{ad} I$ is pure 2-dimensional. Since

$$(R \operatorname{ad} I)|_C^{\mathbf{L}} \cong R\mathcal{H}om_C(I|_C^{\mathbf{L}}, \omega_C) \cong q^n \tag{11.86}$$

is a sheaf, we conclude that $R^1 \operatorname{ad} I$ must be disjoint from C, and thus $\chi(R^1 \operatorname{ad} I) = l|\langle q \rangle|$ for some $l \geq 0$. If $l = 0$, we are done; otherwise, we deduce that $\chi(\operatorname{ad} I) = 1 - n + l|\langle q \rangle| > 1$, contradicting Corollary 11.36. □

Remark. In other words, $I \mapsto \operatorname{ad} I$ gives an isomorphism between the two Calogero-Moser spaces.

Now, consider

$$\kappa_q(\theta I(-D)) \cong R \operatorname{ad} \kappa_q(R \operatorname{ad} I(D + C_m)), \tag{11.87}$$

where D satisfies the hypotheses of Proposition 11.52. Since $\theta^{-1} R \operatorname{ad} I$ is generically an ideal sheaf, Proposition 11.52 applies to show that $\kappa_q(R \operatorname{ad} I(D))(C_m - D)$ is generically an ideal sheaf. But the dual of an ideal sheaf on a *commutative* surface is never a sheaf, and we thus conclude that $\kappa_q(\theta I(-D))$ is never a sheaf, so that we cannot obtain an analogue of Proposition 11.52 in this way.

Another version of the Hilbert scheme of points arises when q is torsion. As we saw above, in the q torsion case, there are exotic instances of simple 0-dimensional sheaves not supported on a power of the anticanonical curve. This suggests that we should investigate the moduli space of such sheaves. The requisite boundedness is trivial in this case, since any 0-dimensional sheaf is automatically both acyclic and globally generated. We thus immediately find that the moduli space of semistable sheaves with numerical invariants $(0, 0, r)$ is a projective scheme for any r. Let M be a 0-dimensional sheaf. If $M|_C \neq 0$, then M is S-equivalent to the direct sum of $M|_C$ and the image of T_M. Since sheaves supported on C are themselves S-equivalent to sums of structure sheaves of points, we thus find that any 0-dimensional sheaf is S-equivalent to a sum of structure sheaves of points on C and 0-dimensional sheaves disjoint from C. A 0-dimensional sheaf disjoint from C only exists if q has finite order, in which case it is supported on a finite subscheme of the center. We then similarly find that such a sheaf is S-equivalent to a sum of sheaves each of which is a simple module over a fiber of the relevant Azumaya algebra. We thus find that the desired moduli space is the disjoint union of locally closed subschemes $\operatorname{Sym}^{r_1}(C) \times \operatorname{Sym}^{r_2}(X_{\phi_* \circ \rho; 1; C'} \setminus C')$, where $\phi : C \to C'$ is the isogeny corresponding to q and $r_1 + r_2 |\langle q \rangle| = r$.

It remains to understand how these pieces are related. For this, it suffices to understand the case $|\langle q \rangle| = r$, so that the moduli space is the disjoint union of $\operatorname{Sym}^r(C)$ and $X_{\phi_* \circ \rho; 1; C'} \setminus C'$. Let Y be the Zariski closure of the latter in the moduli space. Since the generic sheaf classified by Y satisfies $\theta M \cong M$, it follows that *any* sheaf corresponding to a point of Y must at least be S-equivalent to its image under θ. It follows that Y meets $\operatorname{Sym}^r(C)$ along the image of C' under the map taking a point to its preimage under ϕ. Moreover, for any point $x \in C$, we may apply Lemma 11.14 to the blowup of X_m in x to obtain a *simple* sheaf in the corresponding S-equivalence class. In particular, we find that Y is covered by an open subset of the *smooth* algebraic space classifying simple 0-dimensional sheaves of Euler characteristic r. It follows that Y is smooth, implying that $Y \cong X_{\phi_* \circ \rho; 1; C'}$. We have thus shown the following.

Proposition 11.54. *Let q be an r-torsion point of C, with corresponding isogeny $\phi : C \to C'$. Then for any ρ, the moduli space of semistable 0-dimensional sheaves of Euler characteristic r is isomorphic to the union of $\operatorname{Sym}^r(C)$ and $X_{\phi_* \circ \rho; 1; C'}$ with C' identified with its image under $\phi^* : C' \to \operatorname{Sym}^r(C)$.*

Remark. Note that since the Azumaya algebra is nontrivial, this is not a fine moduli space; the Azumaya algebra is in itself the Brauer obstruction to having a universal family.

Remark. Of course, this gluing applies more generally to determine how $C^{r_1} \times X^{r_2}_{\phi_* \circ \rho; 1; C'}$ maps to the moduli space for larger r.

The most striking thing about the above result is not the result itself, but that the argument and the conclusion are both quite similar to that of [47, Thm. 7.1], a theorem about 1-*dimensional* sheaves on *commutative* surfaces. We will explain this in Section 12 below.

11.3 Semistable 1-dimensional sheaves

Since difference equations correspond to 1-dimensional sheaves, it is natural to spend particular effort on the corresponding moduli spaces.

Lemma 11.55. *Let D_a be an ample divisor class, D an effective divisor class, and l an integer. Then there is an integer n such that for any D_a-semistable sheaf M with numerical invariants $(0, D, l)$, $M(nD_a)$ is globally generated. Moreover, this bound can be chosen uniformly over all $X_{\rho;q;C}$ such that D_a is ample.*

Proof. We may feel free to replace D_a by any multiple of D_a, and may thus ensure that $D \cdot D_a < D_a^2$. We may also feel free to start by twisting M by any multiple of D_a, and may thus arrange to have $\chi(M) > 0$. Note that semistability ensures that for any quotient M' of M, $\chi(M'(nD_a)) > 0$ for all $n \geq 0$.

Now, for $0 \leq n$, let M_n denote the image of the natural map

$$\mathcal{O}_X(-nD_a) \otimes_k \mathrm{Hom}(\mathcal{O}_X(-nD_a), M) \to M. \tag{11.88}$$

Since $M_n(nD_a)$ is globally generated, $M_n((n+1)D_a)$ is again globally generated, but now also acyclic by Lemma 11.41. It follows immediately that M_{n+1} contains the saturation $\tilde{M}_n$ of M_n. Moreover, since

$$\chi((M/\tilde{M}_n)((n+1)D_a)) > 0, \tag{11.89}$$

$(M/\tilde{M}_n)((n+1)D_a)$ has a global section, implying that M_{n+1} is strictly larger than $\tilde{M}_n$. In particular, $c_1(M_{n+1}) > c_1(M)$, and thus the sequence must reach M after at most $D \cdot D_a$ steps. □

Since the condition on D_a and the length of the process depend only on $D \cdot D_a$, we conclude the following.

Corollary 11.56. *Let D_a be an ample divisor class. For each integer $r \geq 0$, any semistable 1-dimensional sheaf M with $c_1(M) \cdot D_a = r$ and $\chi(M) > r^2(r+1)$ is acyclic and globally generated.*

Proof. If $c = \lfloor \frac{r}{D_a^2} + 1 \rfloor$, then $(cD_a)^2 > c_1(M) \cdot (cD_a)$, and thus the above argument applies and tells us that if M_0 is a semistable sheaf with $c_1(M_0) = c_1(M)$ and $\chi(M_0) > 0$, then $M_0(rcD_a)$ is acyclic and globally generated. Thus M will be acyclic and globally generated as long as

$$0 < \chi(M(-rcD_a)) = \chi(M) - r^2 c. \tag{11.90}$$

Since $c \leq r + 1$, the claim holds. □

This will allow us to bound the number of global sections of a semistable sheaf, but we will first need to know something about the generic behavior of cokernels of morphisms of line bundles.

Proposition 11.57. *Let D be a universally nef divisor class such that $D \cdot C_m \geq 2$, and let Q be the family of quotients $\mathcal{O}_X/\mathcal{O}_X(-D)$. Then for any pure 1-dimensional sheaf M, there is a fiber Q_v such that $\mathrm{Hom}(Q_v, M) = 0$.*

Proof. Otherwise, there are nonzero morphisms from the generic fiber to M. Let M' be the image of such a morphism, and note that since the generic fiber of Q is transverse to C, so is M'. If M is not transverse to C, we may thus replace it by the kernel of the map $M \to M|_C$, reducing to the transverse case. Since the restriction of the generic fiber of Q to C then has disjoint support from $M|_C$, we conclude that M' must be disjoint from C. Since M' is also globally generated, we may replace M by its maximal globally generated subsheaf which is disjoint from C.

Now, since every fiber still has morphisms to M, this in particular applies to those for which the map $\mathcal{O}_X(-D) \to \mathcal{O}_X$ factors through $T_{\theta^{-1}\mathcal{O}_X(-D)}$. But this implies that the generic quotient $\mathcal{O}_X/\mathcal{O}_X(C_m - D)$ also maps to M. For $m \geq 8$, this lets us reduce to the case $D \cdot e_m = 0$. Since disjointness from C implies that M is α_{m*}-acyclic with pure 1-dimensional direct image, we can reduce to X_{m-1} in such cases. (Note that one must repeat the reduction making M disjoint from C.)

For $m = 7$, the same reduction applies as long as $X_6(D) \neq 0$, since then we still have $(D - C_7) \cdot C_7 \geq 2$. The remaining case is $D = aC_7$ for $a \geq 2$, and in that case the cokernel of T^a has no morphism to M. There is then no difficulty for $m < 7$ until one reaches X_0, X_0' or X_{-1}. Since there are no pure 1-dimensional sheaves disjoint from C in the latter two cases, only X_0 remains.

Here we can reduce as above to $D = df$ or $D = s + df$. For the kernel of a morphism $Q_v \to M$ to be effective, the image must have Chern class $s - f$, and we must have $D = s + df$. Then the kernel has Chern class $(d+1)f$, and is thus an extension of sheaves of Chern class f by Lemma 11.40. But this implies that the points in the support of $Q_v|_C$ can be paired so that their products are in $q^{\mathbb{Z}}\eta$, and this does not happen generically. □

Lemma 11.58. *If M is a semistable 1-dimensional sheaf with $c_1(M) \cdot D_a = r$, then*

$$h^0(M) \leq \max((r+1)^3 + \chi(M), 0). \tag{11.91}$$

Proof. Using $_^D$, we may instead prove that

$$h^1(M) \leq \max((r+1)^3 - \chi(M), 0). \tag{11.92}$$

If $\chi(M) > r^2(r+1)$, this follows from Corollary 11.56. Otherwise, let $c = \lceil (r+1)^2 - \frac{\chi(M)}{r} \rceil$, so that $\chi(M(cD_a)) \geq r(r+1)^2 > r^2(r+1)$, and thus $M(cD_a)$ is acyclic. It follows from the previous Lemma that there is a quotient $N \cong \mathcal{O}_X/\mathcal{O}_X(-cD_a)$ such that the corresponding map

$$\mathrm{Hom}(\mathcal{O}_X, M) \to \mathrm{Hom}(\mathcal{O}_X(-cD_a), M) \tag{11.93}$$

is injective. Thus $\mathrm{Hom}(N, M) = 0$ while $\mathrm{Ext}^2(N, M) \cong \mathrm{Ext}^2(\mathcal{O}_X, M) = 0$. We deduce that

$$h^1(M) \leq \dim \mathrm{Ext}^1(N, M) = -\chi(N, M) = cD_a \cdot M = cr < (r+1)^3 - \chi(M). \tag{11.94}$$

□

This may be used in place of the Le Potier-Simpson estimate (see [26, Thm. 3.3.1] for the specific version we adapted) to show that sheaves with sufficiently small slope cannot be destabilizing for the GIT quotient. Following the standard argument from the commutative case [26, §4.4] gives us the following, where we call a pair $(D, l) \in \mathrm{Pic}(X) \times \mathbb{Z}$ *primitive* if there is no integer $r > 1$ such that $D \in r\,\mathrm{Pic}(X)$, $l \in r\mathbb{Z}$.

Theorem 11.59. *Let D be an effective divisor class, and let D_a be an ample divisor class. Then the moduli problem of classifying S-equivalence classes of D_a-semistable sheaves M with numerical invariants $(0, D, l)$ is represented by a projective scheme $M_{D,l}$. The stable locus of $M_{D,l}$ is a Poisson subscheme, which is smooth if $D \cdot C_m > 0$, as is the Zariski closure of the locus with $M|_C = 0$ if $D \cdot C_m = 0$. If (D, l) is primitive, the stable locus admits a universal family.*

Proof. Projectivity follows as discussed above. That the stable locus is Poisson follows from the fact that stable sheaves are simple. If $D \cdot C_m > 0$, then stability implies $\mathrm{Hom}(M, \theta M) = 0$ (since the slope drops) and thus deformations are unobstructed. Similarly, if $D \cdot C_m = 0$, then M and

θM are stable sheaves of the same slope, and thus any map is an isomorphism. Semicontinuity then implies that $\dim \mathrm{Ext}^2(M,M) = \dim \mathrm{Hom}(M,\theta M) = 1$ on the Zariski closure of the given symplectic leaf, and thus $\dim \mathrm{Ext}^1(M,M)$ is constant, so that the variety is generically reduced with constant tangent space dimension.

Moreover, the stable locus admits a universal family as an algebraic space, which means that as a quasiprojective variety, it has a universal family étale locally. Let M denote this family. The obstruction to M descending to a family on the quasiprojective moduli space lies in the Brauer group of the moduli space, and the usual argument shows that the element in the Brauer group is annihilated by $\chi(\mathcal{O}_X(D'), M)$ for all D'. Since (D, l) is primitive, the set of numbers $\chi(\mathcal{O}_X(D'), M)$ has gcd 1, and thus the obstruction vanishes. □

Remark 1. Note that if (D, l) is primitive and $l \neq 0$, we can always choose the ample divisor so that every semistable sheaf is stable, by taking D_a so that $D \cdot D_a$ is prime to l. In fact, given any choice of ample bundle, we can modify it so that not only is every semistable sheaf stable, but all of the sheaves which were stable for the original bundle remain stable. To do this, simply observe that there are only finitely many possible divisor classes of subsheaves of a sheaf with Chern class D, and thus a discrete set of possible slopes of subsheaves. In particular, there is a lower bound on a nonzero difference between the slope of a subsheaf and $\mu = l/D \cdot D_a$. Thus if we could replace D_a by $D_a + \epsilon D_0$ for ϵ sufficiently small, the slopes on either side of μ would remain sufficiently small. Replacing D_a by $nD_a + D_0$ for $n \gg 0$ thus has a similar effect (essentially using D_0 to break ties between sheaves of the same slope). The resulting smooth projective Poisson moduli space has a natural morphism to the original semistable moduli space, which thus inherits a Poisson structure. (We can do something similar if $l = 0$ if we first twist by D_a before changing the ample bundle.)

Remark 2. There are imprimitive examples, even in the commutative case, for which the Brauer obstruction to the existence of a universal family is nontrivial, see the remark following [47, Prop. 6.3] or Proposition 12.8 below.

Remark 3. Note that any disjoint-from-C case with a universal family gives rise to a Lax pair, in the following way. Let the difference equation corresponding to the universal family (on the main symplectic leaf) be

$$v(qz) = A(z; \mu; \rho; q; p)v(z), \tag{11.95}$$

where μ is a point of the moduli space. If we twist the sheaf by a line bundle, then (as we will see below) the new equation has the form

$$\hat{v}(qz) = C(qz; \mu; \rho; q; p)A(z; \mu; \rho; q; p)C(z; \mu; \rho; q; p)^{-1}\hat{v}(z) \tag{11.96}$$

where C is a meromorphic map between the corresponding vector bundles. The resulting sheaf (assuming it is still stable, e.g., if the sheaf is irreducible) is of the form classified by a related moduli space, and is thus equivalent to the equation

$$v(qz) = A(z; \phi(\mu); \tilde{\rho}; q; p)v(z) \tag{11.97}$$

where $\tilde{\rho}$ are the new parameters and ϕ is a morphism between the stable loci of the two moduli spaces. But then there is an isomorphism C_0 between the respective line bundles identifying the two equations. Absorbing this into C produces a pair of equations

$$\begin{aligned} v(qz; \mu; \rho; q; p) &= A(z; \mu; \rho; q; p)v(z; \mu; \rho; q; p) \\ v(z; \phi(\mu); \tilde{\rho}; q; p) &= C(z; \mu; \rho; q; p)v(z; \mu; \rho; q; p) \end{aligned}$$

satisfying the consistency condition

$$A(z;\phi(\mu);\tilde{\rho};q;p)C(z;\mu;\rho;q;p) = C(qz;\mu;\rho;q;p)A(z;\mu;\rho;q;p), \tag{11.98}$$

and thus giving the desired Lax pair.

This fails, of course, when there is no universal family on the moduli space, though there are possible approaches to work around this issue. For instance, as suggested by [41], one could choose a trivialization of $\rho_* M$ at some auxiliary point u not in the orbit of any singularities. We can then carry this trivialization through the twist by any divisor class in $\langle f, e_1, \dots, e_m\rangle$, and thus that subgroup of $\mathrm{Pic}(X)$ acts on the new moduli space, a PGL_n-bundle over the true moduli space. Moreover, this new moduli space now has a universal family, and thus one obtains a Lax pair as above. (If one also chooses a trivialization of $\rho_* M(-s)$, one can make all of $\mathrm{Pic}(X)$ act, though now $\mathrm{Pic}(X)$ acts nontrivially on the auxiliary base points.)

We should note here that we have not actually shown that the stable locus is nonempty! In general, it is unclear how to do this (especially since there are certainly cases in which there are no stable sheaves), but there are a few cases in which we have particularly nice constructions. In each case, we obtain an explicit birational map from one component of the stable locus to a corresponding commutative moduli space; in addition, in certain cases, there are natural birational maps between the cases.

The first case is $l = 0$. Since such sheaves have $\chi = 0$, it is natural to consider the case $M \in \mathcal{N}_{\rho;q;C}$ and hope that $\kappa_q M$ will generically be a sheaf. As in the proof of Proposition 11.52, it will be convenient to work backwards instead. Thus let M be a generic semistable sheaf on $X_{\rho;1;C}$ with invariants $(0, D, 0)$. We assume in addition that the linear system $|D|$ has an integral representative (and that M is generic in the corresponding component), since otherwise we have no hope of constructing stable sheaves. In that case, D is represented by an integral curve of genus g, and M is an ineffective line bundle of degree $g-1$ on such a curve. In particular, it follows that M is the minimal lift of a sheaf on X_0. Its image M_0 on X_0 is still in the kernel of $R\Gamma$, and thus (arguing as in the proof of Lemma 11.10) has a presentation

$$0 \to \pi^* W(-s) \to \mathcal{O}_X(-f)^{c_1(M)\cdot f} \to M_0 \to 0 \tag{11.99}$$

where W is a vector bundle on $\mathbb{P}^1$ of rank $c_1(M)\cdot f$. Moreover, this morphism is injective on C, since M_0 is transverse to C. Applying κ_q^{-1} gives essentially the same presentation, and since the morphism is still injective on C, the cokernel will still be a pure 1-dimensional sheaf transverse to C. Since κ_q commutes with the minimal lift, it follows that $\kappa_q^{-1} M$ is a pure 1-dimensional sheaf disjoint from C with the desired Chern class. Moreover, since $\kappa_q^{-1} M$ has no global sections, it follows that any subsheaf of $\kappa_q^{-1} M$ has nonpositive Euler characteristic, and thus $\kappa_q^{-1} M$ is semistable.

Remark. The condition on D is quite weak; per [47, Thm. 4.8] (based on [23, Thm. III.1(c)]), the only universally nef divisor classes disjoint from C_m which are not generically integral are those of the form rC_8 with $\rho(C_8)$ torsion of order strictly dividing r, and of the form $rC_8 + e_8 - e_9$ with $\rho(e_8 - e_9) = 1$, and in neither case are there strictly stable sheaves of Euler characteristic 0.

The significance of the above presentation is that it corresponds directly (after twisting by f) to a symmetric elliptic difference equation of the form

$$v(qz) = A(z)v(z) \quad v(q\eta/z) = v(z) \tag{11.100}$$

with $A(z)$ a matrix (rather than a morphism of vector bundles) satisfying $A(\eta/z)A(z) = 1$, and with singularities depending on the points being blown up and the divisor class D. In particular, the

conditions on A are independent of q, which is why $\kappa_{q'}^{-1}\kappa_q$ simply changes q to q' in the difference equation.

The next case to consider is $l = -1$. The key idea in this case is that if such a sheaf has no global sections (which one would expect, by analogy with the standard commutative heuristic, to be the complement of a codimension $1 - \chi(M) = 2$ condition), then there is a universal extension

$$0 \to \theta\mathcal{O}_X \to I \to M \to 0, \tag{11.101}$$

where I has invariants $(1, D - C_m, 0)$ and $R\Gamma I = 0$. In particular, $I(C_m - D)$ corresponds to a point of $\mathrm{Hilb}^{D\cdot(D-C_m)/2+1}(X_{\rho;q;C}(C_m - D))$. Since $(D - C_m)\cdot C_m > 0$ (unless $D = rC_m$ for some r), we understand what the generic such sheaf looks like, suggesting that we should construct M by reversing this construction.

It turns out that the resulting construction works as long as either $\rho(D) = q$ or $D \cdot C_m > 0$. Let Z be a generic $D \cdot (D - C_m)/2 + 1$-point subscheme in $X_{\rho;1;C}$, and consider the sheaf

$$I := \kappa_q^{-1} I_Z(D - C_m) \tag{11.102}$$

on $X_{\rho;q;C}$. Then $R\,\mathrm{Hom}(\theta\mathcal{O}_X, \theta I)$=0, and thus

$$\mathrm{Hom}(\theta\mathcal{O}_X, I) \cong R\,\mathrm{Hom}(\theta\mathcal{O}_X, I|_C) \cong R\Gamma(C; (\theta^{-1}I)|_C) \cong R\Gamma(C; q^{D\cdot C_m - 1}\rho(D)) \tag{11.103}$$

In particular, we obtain a $\mathbb{P}^{\max(0, D\cdot C_m - 1)}$ of such maps up to scalar multiples. The usual argument says that any such nonzero morphism is injective, and thus we have a short exact sequence

$$0 \to \theta\mathcal{O}_X \to I \to M \to 0. \tag{11.104}$$

Now, M certainly has the correct invariants, so it remains only to show that it is stable. Since $R\Gamma(I) = 0$, we find $\Gamma(M) = 0$, and thus no subsheaf has a global section. Thus the only way a subsheaf N could be destabilizing would be if it also had no H^1. But in that case, the inclusion map $N \to M$ would lift to a map $N \to I$, contradicting pure 2-dimensionality of I. We thus obtain a rational map from a $\mathbb{P}^{\max(0, D\cdot C_m - 1)}$-bundle over $\mathrm{Hilb}^{D\cdot(D-C_m)/2+1}(X_{\rho;1;C})$ to the moduli space of stable sheaves of Chern class D and Euler characteristic -1. (Note that both spaces have dimension $D^2 + 1$ if $D \cdot C_m > 0$ or $D^2 + 2$ if $\rho(D) = q$.)

The case $l = 1$ is then straightforward: if M is a stable 1-dimensional sheaf of Euler characteristic -1, then $R^1 \operatorname{ad} M$ is a stable 1-dimensional sheaf of Euler characteristic 1, and vice versa. If q is not torsion (or the arithmetic genus of D is sufficiently small compared to the order of q), then we can dualize the description of M in terms of I, obtaining a short exact sequence

$$0 \to \operatorname{ad} I \to \mathcal{O}_X \to R^1 \operatorname{ad} M \to 0. \tag{11.105}$$

In other words, the sheaves with $l = 1$ we obtain are those that are generated by their unique global section. This, too, has a natural interpretation in terms of difference equations: these are the "straight-line" difference equations. Indeed, we generically have a unique global section

$$\mathrm{Hom}(\mathcal{O}_X(-D - gf), \operatorname{ad} I) \tag{11.106}$$

which combines to give an expression for $R^1 \operatorname{ad} M$ as a quotient of some quotient $\mathcal{O}_X/\mathcal{O}_X(-D-gf)$. In other words, there is an operator $\mathcal{D} \in \hat{\mathcal{S}}'_{\rho;q;C}(0, D + gf)$ such that $R^1 \operatorname{ad} M$ corresponds to the equation $\mathcal{D}v = 0$. (The kernel of the map to $R^1 \operatorname{ad} M$ has no maps or Ext^1 to $\mathcal{M}er$, see below; the only effect is to introduce g apparent singularities.)

Note that in the case $D \cdot C_m = 0$, the construction for $l = 1$ gives an acyclic globally generated sheaf with $M \cong \theta M$. But then the map from M to I is the same (anti-)Poisson map we used to prove the Jacobi identity. Since κ_q is also Poisson, we obtain a birational symplectic map between the moduli space classifying M and the appropriate commutative Hilbert scheme.

Similar ideas give us the following.

Proposition 11.60. *The generic sheaf in any irreducible component of the subscheme of $M_{D,l}$ classifying irreducible sheaves disjoint from C satisfies either $H^0(M) = 0$ or $H^1(M) = 0$.*

Proof. Using the adjoint, it suffices to prove that $H^0(M) = 0$ for the generic such sheaf when $l \leq 0$. By tangent space considerations, the subscheme of irreducible sheaves has dimension $D^2 + 2$ everywhere, and thus it will suffice to prove that the locally closed subscheme M_n where $h^0(M) = n$ has strictly smaller dimension for $n > 0$. The corresponding moduli space of (surjective, by irreducibility) maps $\mathcal{O}_X \to M$ is a $\mathbb{P}^{n-1}$-bundle over M_n, so for $n > 0$ has dimension $n-1+\dim M_n$. On the other hand, the kernel of such a map has the form $I(-D)$ for $I \in \mathrm{Hilb}^{D^2/2-l}$, and one has a short exact sequence

$$0 \to \mathrm{Hom}(O_X, O_X) \to \mathrm{Hom}(I(-D), O_X) \to \mathrm{Ext}^1(M, O_X) \to 0 \tag{11.107}$$

with $\dim \mathrm{Ext}^1(M, O_X) = \dim H^1(M) = n + l$. Thus the moduli space of maps $\mathcal{O}_X \to M$ can also be viewed as a $\mathbb{P}^{n+l}$-bundle over a subscheme of $\mathrm{Hilb}^{D^2/2-l}$, so has dimension at most $D^2 + n - l$. Comparing gives

$$\dim M_n \leq D^2 + 2 - (l + 1) < D^2 + 2 \tag{11.108}$$

as required. □

Remark. Note that unless $D = C_m = C_8$, an irreducible sheaf with $c_1(M) = D$ must be disjoint from C, since otherwise $M|_C^{\mathbf{L}}$ would be a 1-dimensional subsheaf.

This implies that the above construction of 1-dimensional sheaves from sheaves in the Hilbert scheme is generically invertible for irreducible sheaves disjoint from C_0, giving the following.

Corollary 11.61. *The subscheme of $M_{D,1}$ classifying irreducible sheaves disjoint from C is either empty or birational to $\mathrm{Hilb}^{1+D^2/2}(X(D))$. The subscheme of $M_{D,-1}$ classifying irreducible sheaves disjoint from C is either empty or birational to $\mathrm{Hilb}^{1+D^2/2}(X(-D))$.*

Remark. One can show that, with a few exceptions, the above construction generically produces irreducible sheaves, and thus rules out the "empty" case above. The idea (we omit the details) is that the inductive construction of Proposition 11.52 in most case produces an extension of two irreducible sheaves, and thus one need simply show that there are sheaves without a component of either of those Chern classes (either as a consequence of transversality or by a dimension computation). The few possible universally nef exceptions (arising either from failures in base cases or the fact that blowing up can introduce components) are the ones one would expect from the commutative setting (see [47, §4.3], also [23]):

(1) $D = rf$, $r > 1$

(2) For some $1 \leq i \leq m$, $D \cdot e_i = 0$ and $\rho(D) \in x_i q^{\mathbb{Z}}$

(3) $D = rC_8 + e_8$, $\rho(C_8) \in q^{\mathbb{Z}}$

(4) $D = rC_8$, $q\rho(C_8)^r = 1$, $\rho(C_8)^{r'} = 1$ for some $r' < r$.

(5) $D = rC_8 + e_8 - e_9$, $q\rho(D) = 1$, $x_8/x_9 \in q^{\mathbb{Z}}$

The given conditions are where the simplest argument fails; it is conceivable that there may be some instances of (2–5) where the stable sheaves from the above construction are generically irreducible. Note that (4) and (5) are the only cases which can be disjoint from C_m, and (4) forces q to be torsion.

If $c_1(M) \cdot e_m = 1$, then $\chi(M(e_m)) = \chi(M) + 1$, and thus twisting by e_m relates sheaves with $\chi = 0$ to sheaves with $\chi = 1$. It turns out that one can give an explicit description of this action on generic sheaves, in terms of the respective commutative data. It is somewhat easier to describe this in terms of sheaves on X_{m-1}, bearing mind that M is generically a minimal lift.

In particular, given generic M on X_{m-1} with invariants $(0, D, 0)$, $D \cdot C_{m-1} = 1$, coming from $\kappa_q^{-1}(\mathcal{L})$ with $\mathcal{L}$ an ineffective degree $g-1$ line bundle on an irreducible curve in $|D|$, we want to understand the point in Hilb^n corresponding to

$$M' = \theta \alpha_{m*} \theta^{-1} \alpha_m^{*!} M. \tag{11.109}$$

We recall from Corollary 11.19 that this is the universal extension

$$0 \to M \to M' \to \mathcal{O}_{qx_m} \to 0; \tag{11.110}$$

here we use the fact that $\alpha_m^{*!} M$ is disjoint from C to compute $\dim \mathrm{Ext}^1(\mathcal{O}_{qx_m}, M) = 1$. Similarly, we can reconstruct M from M' as the kernel of the universal map to $\mathcal{O}_{qx_m}$.

To relate this to the commutative Hilbert scheme, we need to consider the kernel K of the unique global section of M'. An easy snake lemma calculation gives a short exact sequence

$$0 \to K \to I_{qx_m} \to M \to 0 \tag{11.111}$$

Each sheaf in this short exact sequence is in the kernel of $R\Gamma$, so we may apply κ_q to obtain a distinguished triangle

$$\kappa_q K \to I_{x_m} \to \mathcal{L} \to \tag{11.112}$$

To relate this to the Hilbert scheme, we observe that both I_{x_m} and $\mathcal{L}$ have a unique extension by $\mathcal{O}_{x_m}$ (which is the point of intersection of the support of $\mathcal{L}$ with C), and thus the complex

$$I_{x_m} \longrightarrow \mathcal{L} \tag{11.113}$$

is quasi-isomorphic to the complex

$$\mathcal{O}_X \longrightarrow \mathcal{L}(x_m). \tag{11.114}$$

The dual of this complex is, up to twisting, the ideal sheaf of the unique effective representative of the divisor class of $\mathcal{L}(x_m)$ on its support.

We thus obtain the following description: Starting with a line bundle $\mathcal{L}$ on a curve C' in $|D|$, we twist up by the unique point of intersection, and view the corresponding effective divisor on C' as a subscheme of X. The reverse operation is equally straightforward: given a $g(D)$-point subscheme of X, there is generically a unique curve in $|D|$ passing through that subscheme, producing a line bundle of degree g on that curve. Twisting down by the point of intersection gives a line bundle of degree $g-1$, which will be generically ineffective.

This procedure is most interesting when we also have $D \cdot e_{m-1} = 1$, as we can then take the commutator of the above operation with reflection in $e_{m-1} - e_m$ to twist by $e_{m-1} - e_m$. This should be compared with the description of elliptic Painlevé in [29] (to which it is related by a derived equivalence, see below); it is also a direct generalization of [39, Thm. 2].

Of course, we can similarly twist between $\chi = -1$ and $\chi = 0$; since minimal lifting and κ_q both commute with R^1 ad, this is just the conjugate of the above operation by duality. It would also be of interest to understand twisting between $\chi = -1$ and $\chi = 1$ in the case $D \cdot e_m = 2$.

11.4 Lax pairs for elliptic Painlevé

When $D \cdot C_m = D^2 = 0$, the above moduli space is a quasiprojective surface, which raises the natural question of whether we can identify the components. We may as well assume that D is in the fundamental chamber, as other cases will be related by a sequence of reflections (or reflection functors). If $D \cdot e_m < 0$, then D cannot be disjoint from C_m, and thus we reduce to the case that D is universally nef. It is then straightforward to see that a universally nef divisor satisfying $D \cdot C_m = D^2 = 0$ must be a multiple of C_8. In other words, the 2-dimensional cases are essentially just the moduli spaces classifying sheaves on X_8 with $c_1(M) = rC_8$, $\chi(M) = d$.

For convenience in notation, we fix $\eta, x_0, \dots, x_7$ and reparametrize by $x_8 = (\eta^2 x_0^3 / x_1 \cdots x_7 z)$ to obtain a surface $X_{z,q}$. As we have noted above, the parameters of $X_{\rho;q;C}$ can be recovered from the restrictions of various classes of $K_0(X_{\rho;q;C})$ to C; the parameters z and q in particular corresponding to restrictions of sheaves supported on C:

$$q = c_1(\mathcal{O}_x|_C^{\mathbf{L}}), \quad z = c_1(\mathcal{O}_C|_C^{\mathbf{L}}). \tag{11.115}$$

Note in particular that if $c_1(M) = rC_8$, $\chi(M) = d$, then $c_1(M|_C^{\mathbf{L}}) = q^d z^r$, so for such a sheaf to be disjoint from C requires that $q^d z^r = 1$. If r and d are relatively prime, this is equivalent to asking that $z = w^{-d}$, $q = w^r$ for some w.

Theorem 11.62. *Suppose $r > 0$ and d are relatively prime integers, and let D_a be an ample divisor class such that any semistable 1-dimensional sheaf on X_{w^{-d},w^r} with $c_1(M) = rC_8$, $\chi(M) = d$ is actually stable. Then the corresponding semistable moduli space is isomorphic to $X_{w,1}$.*

Proof. Since the tangent space has dimension

$$\dim \operatorname{Ext}^1(M, M) = \dim \operatorname{Hom}(M, M) + \dim \operatorname{Hom}(M, \theta M) = 2 \tag{11.116}$$

at the point corresponding to M, we find that every component is either nonreduced or smooth, and must be the latter if it contains a sheaf disjoint from C. Moreover, the closed subscheme corresponding to sheaves supported on C is just the moduli space of rank r degree d vector bundles on C, so is naturally isomorphic to $\operatorname{Pic}^d(C)$, and is in particular connected and 1-dimensional.

Fix a component Y containing sheaves disjoint from C (a smooth projective surface), and let F be a corresponding universal family (which exists since r and d are relatively prime); that such a component exists will be shown in Corollary 12.4 below. Just as in the commutative setting, this gives rise to a Fourier-Mukai functor $\Phi_F : D^b \operatorname{coh} Y \to D^b \operatorname{coh} X_{w^{-d},w^r}$. Indeed, we may view F as a sheaf on the (flat) base change of X_{w^{-d},w^r} to Y, and thus we can tensor F with any complex on Y, then take the derived direct image to X_{w^{-d},w^r}. We claim that Φ_F is actually an equivalence between the two derived categories. By [13, Thm. 2.4] (using the fact that point sheaves are a spanning class for $D^b \operatorname{coh} Y$), it suffices to prove the following:

(1) Φ_F has both a left and a right adjoint.

(2) For all $x \in Y$, $\Phi_F(\mathcal{O}_x \otimes \omega_Y) \cong \theta \Phi_F(\mathcal{O}_x)$.

(3) For any two points $x, y \in Y$ and $p \in \{0, 1, 2\}$, the natural morphism

$$\operatorname{Ext}^p(\mathcal{O}_x, \mathcal{O}_y) \to \operatorname{Ext}^p(\Phi_F(\mathcal{O}_x), \Phi_F(\mathcal{O}_y)) \tag{11.117}$$

is a bijection.

(4) The category $D^b \operatorname{coh} X_{w^{-d},w^r}$ is indecomposable.

For (1), the right adjoint applied to N is $R\mathcal{H}om_Y(F, N_Y)$, just as for commutative Fourier-Mukai functors, and conjugating by the respective Serre functors gives a left adjoint.

For (2) and (3), note that $\Phi_F(\mathcal{O}_x)$ is just the fiber at x of the universal family. For (2) we thus reduce to the fact that every sheaf in the family satisfies $\theta M \cong M$. For (3) with $x \neq y$, we have two nonisomorphic sheaves of the given kind, and stability prevents there being a morphism in either direction. But then Euler characteristic considerations imply that $R\operatorname{Hom}(\Phi(\mathcal{O}_x), \Phi(\mathcal{O}_y)) = R\operatorname{Hom}(\mathcal{O}_x, \mathcal{O}_y) = 0$ as required. Similarly, $R\operatorname{Hom}(\Phi(\mathcal{O}_x), \Phi(\mathcal{O}_x))$ has the same Betti numbers as $R\operatorname{Hom}(\mathcal{O}_x, \mathcal{O}_x)$, since $\operatorname{Hom}(\Phi_F(\mathcal{O}_x), \theta\Phi_F(\mathcal{O}_x)) \cong \operatorname{Hom}(\Phi_F(\mathcal{O}_x), \Phi_F(\mathcal{O}_x)) = k$.

Now, the map $\operatorname{Hom}(\mathcal{O}_x, \mathcal{O}_x) \to \operatorname{Hom}(\Phi_F(\mathcal{O}_x), \Phi_F(\mathcal{O}_x))$ certainly takes the identity to the identity; since both sides are spanned by the identity, this is an isomorphism. Next, the map $\operatorname{Ext}^1(\mathcal{O}_x, \mathcal{O}_x) \to \operatorname{Ext}^1(\Phi_F(\mathcal{O}_x), \Phi_F(\mathcal{O}_x))$ is just the usual (Kodaira-Spencer) identification of the tangent space to the moduli space with the space of infinitesimal deformations. Finally, the map $R\operatorname{Hom}(\mathcal{O}_x, \mathcal{O}_x) \to R\operatorname{Hom}(\Phi_F(\mathcal{O}_x), \Phi_F(\mathcal{O}_x))$ respects the Yoneda pairing, and Serre duality implies that the image of $\operatorname{Ext}^1(\mathcal{O}_x, \mathcal{O}_x) \otimes \operatorname{Ext}^1(\mathcal{O}_x, \mathcal{O}_x)$ spans $\operatorname{Ext}^2(\mathcal{O}_x, \mathcal{O}_x)$.

Finally, for indecomposability, we simply note that since $R\operatorname{Hom}(\mathcal{O}_X(-D), \mathcal{O}_X(-D)) = k$, each line bundle must actually be in a component, and two line bundles differing by an ample divisor class must be in the same component. But this implies that there is a component containing all line bundles, and line bundles generate.

We have thus shown that Φ_F is a derived equivalence. Now, the inverse of Φ_F transports the natural transformation $T : \theta \to \mathrm{id}$ to a natural transformation $T : _ \otimes \omega_X \to \mathrm{id}$, so that Y has an anticanonical curve C'. Moreover, we find that Φ_F restricts to an equivalence $D^b \operatorname{coh} C' \cong D^b \operatorname{coh} C$. Since the derived category of a smooth curve determines the curve, we conclude that $C' \cong C$. From the classification of Poisson surfaces [45], any Poisson surface with smooth, connected anticanonical curve is rational. Moreover, Φ_F identifies $K_0(Y)$ and $K_0(X_{w^{-d},w^r})$, and thus $K_Y^2 = 0$.

Thus Y can be expressed as the blowup of F_1 in 8 points of a suitably embedded copy of C; we fix such an expression to obtain a basis of $K_0(Y)$. The action of Φ_F respects restriction to C (up to an autoequivalence of C), so respects both the image and the kernel of the composition

$$R_C : K_0(X) \to K_0(C) \to K_0(X), \tag{11.118}$$

and similarly for Y. Since $R_C^2 = 0$, it is natural to consider $\ker(R_C)/\operatorname{im}(R_C)$. The natural pairing induced on this subquotient identifies it with the lattice Λ_{E_8}, and the action of $W(E_9)$ on either side acts as $W(E_8)$ on this lattice. In particular, there exists a description of Y as a blowup of F_1 with the property that Φ_F identifies the corresponding bases of this subquotient. Now, consider how Φ_F acts on the kernel of R_C. For each simple root r_i of E_8, there is a corresponding class $[\mathcal{O}_{r_i}(-1)]$ of $K_0(X)$ (i.e., the class with numerical invariants $(0, r_i, 0)$), and Φ_F^{-1} of this class will have rank 0 and $c_1 \in r_i + \langle C_8 \rangle$. The action of $\Lambda_{E_8} \subset W(E_9)$ lets us make the Chern class of each $\Phi_F^{-1}[\mathcal{O}_{r_i}(-1)]$ equal to r_i (again, at the cost of changing how we represent Y as a blowup), and the action of Λ_{E_8} as twisting by line bundles then lets us make the Euler characteristics 0. (Note that this twists F by a line bundle (and thus changes the derived equivalence), but universal families are only defined up to twisting by a line bundle in any event.) In other words, we may assume that our blowdown structure on Y and universal family F are such that $\Phi_F[\mathcal{O}_{r_i}(-1)] = [\mathcal{O}_{r_i}(-1)]$ for any simple root of E_8. Moreover, since Φ_F must take $\operatorname{im}(R_{C'})$ to $\operatorname{im}(R_C)$, there is a corresponding element of $\mathrm{GL}_2(\mathbb{Z})$. Since

$$\Phi_F([\mathcal{O}_x]) = r[\mathcal{O}_C] + d[\mathcal{O}_x], \tag{11.119}$$

we find that

$$\Phi_F([\mathcal{O}_{C'}]) = a[\mathcal{O}_C] + b[\mathcal{O}_x] \tag{11.120}$$

where $ad - br \in \{\pm 1\}$.

In particular, up to the uncertainty in a and b, we have determined how Φ_F acts on $\mathrm{im}(R_{C'})$. Since we have a commutative diagram

$$\begin{array}{ccc} \mathrm{im}(R_{C'}) & \xrightarrow{\Phi(F)} & \mathrm{im}(R_C) \\ \downarrow & & \downarrow \\ K^0(C') & \xrightarrow{\Phi(F)} & K^0(C), \end{array} \tag{11.121}$$

this is enough to compute how $|^{\mathbf{L}}_{C'}$ acts on $\mathrm{im}(R_{C'})$. In particular, if Y has parameters $\tilde{\eta}, \tilde{x}_0, \dots, \tilde{x}_8$, then (identifying $\mathrm{Pic}^0(C')$ with $\mathrm{Pic}^0(C)$)

$$\begin{aligned} \tilde{x}_i/\tilde{x}_{i+1} &= x_i/x_{i+1}, \quad 0 \le i \le 6, \\ \tilde{\eta}/\tilde{x}_1\tilde{x}_2 &= \eta/x_1x_2. \end{aligned}$$

But then we need merely multiply the isomorphism $C \cong C'$ by a translation on C' to also ensure that $\tilde{x}_7 = x_7$, so that $Y \cong X_{z,1}$ for some z. To compute z, we just need to restrict $[\mathcal{O}_{C'}]$ to C'; under Φ_F, this becomes $w^{-ad+br} \in w^{\pm 1}$, and thus either $Y \cong X_{w,1}$ or $Y \cong X_{1/w,1}$. Since these two surfaces are isomorphic, the claim follows.

To see that this is the only component, observe that it follows from the above argument that any component containing sheaves disjoint from C also contains the curve $C' \cong \mathrm{Pic}^d(C)$ of sheaves supported on C. It follows that the moduli space is smooth and connected, so irreducible as required. □

Remark. Note that since there can be multiple inequivalent choices of stability condition, the above derived equivalences are not canonical.

In particular, we see that for any $d/r \in \mathbb{Q} \cup \{\infty\}$ there is an interpretation of $X_{w,1}$ as a moduli space of sheaves. Moreover, any element of $\mathrm{Pic}(X) \rtimes W(E_9)$ induces isomorphisms between the moduli spaces of sheaves and thus between surfaces of the form $X_{w,1}$. In the case $r = 0$, $\mathrm{Pic}(X)$ acts trivially, and thus one has an action of $W(E_9)$ alone, the translation subgroup of which is the elliptic Painlevé equation of [54]. In general, one can transport the action through the derived equivalence to find that the kernel is still isomorphic to $\mathrm{Pic}(X)$, but now with a slightly different embedding. In each case, we find that the complementary subgroup of $\mathrm{Pic}(X) \rtimes \Lambda_{E_8} \cong \mathbb{Z}.\Lambda^2_{E_8}$ acts on the moduli spaces as the elliptic Painlevé equation. In particular, we find that twisting by line bundles (which we will see acts on the difference equations as isomonodromy transformations) acts on the moduli space as $r\Lambda_{E_8} \subset W(E_9)$. As a result, every such moduli problem gives rise to a Lax pair for the elliptic Painlevé equation (decimated by a factor of $|r|$); see the remark following Theorem 11.59, corresponding to a symmetric elliptic difference equation of order $2r$ on a vector bundle of degree d such that $A(x_i)$ has rank r for $1 \le i \le 8$. The Lax pair of [67], which was given in straight-line form, presumably corresponds to the case $r = d = 1$. Similarly, the Lax pair of [50] corresponds to the "hypergeometric solution" in the case $r = 1$, $d = 0$. More precisely, it corresponds to the case in which M has a quotient $\mathcal{O}_{s-f}(n-1)$ for $n \ge 0$, which for $n > 0$ is stable, corresponding to a point on an appropriate -2-curve on the moduli space. (In fact [50] considered the more general case in which $c_1(M) = 2s + (l+2)f - e_1 - \cdots - e_{2l+6}$, $\chi(M) = 0$, and M has a quotient of the form $\mathcal{O}_{s-f}(n-1)$ for $n \ge 0$. For the general case with these numerical invariants, see [41].)

In the case $d = 1$, we can use the Hilbert scheme idea to give a relatively explicit construction of the universal family; more precisely, the Hilbert scheme construction gives a generically stable (in fact, irreducible) family, and the universal family will be obtained by taking stable limits.

Let w be an element of $\mathrm{Pic}^0(C)$, r a positive integer less than $|\langle w\rangle|$, and let I be a pure 2-dimensional sheaf on X_{w^{-1},w^r} with numerical invariants $(1,0,0)$. As we have seen above, the moduli space of such I is naturally identified with $X_{w,1}$, so there certainly exists such an I, and we may further insist that T_I is injective. Now, consider the distinguished triangle

$$R\,\mathrm{Hom}(\theta^l I, \theta\mathcal{O}_X) \to R\,\mathrm{Hom}(\theta^l I, \mathcal{O}_X) \to R\,\mathrm{Hom}_C(\theta^l I|_C, \mathcal{O}_C) \to \tag{11.122}$$

Since $\theta^l I|_C$ is the line bundle with isomorphism class w^{l-r}, we find that

$$R\,\mathrm{Hom}(\theta^l I, \mathcal{O}_X) \cong R\,\mathrm{Hom}(\theta^l I, \theta\mathcal{O}_X) \cong R\,\mathrm{Hom}(\theta^{l-1} I, \mathcal{O}_X) \tag{11.123}$$

for $0 \le l < r$. Since $R\,\mathrm{Hom}(\theta^{-1} I, \mathcal{O}_X) = 0$, we conclude by induction that $R\,\mathrm{Hom}(\theta^{r-1} I, \mathcal{O}_X) = 0$. But then the same distinguished triangle gives

$$R\,\mathrm{Hom}(\theta^r I, \mathcal{O}_X) \cong R\,\mathrm{Hom}_C(\theta^r I|_C, \mathcal{O}_C) \cong R\,\mathrm{Hom}_C(\mathcal{O}_C, \mathcal{O}_C). \tag{11.124}$$

In particular, we find that $\mathrm{Hom}(\theta^r I, \mathcal{O}_X) \cong k$.

Fix a generator of that space. The image is a nonzero subsheaf of $\mathcal{O}_X$, so must have rank 1, and thus the kernel has rank 0, and thus vanishes since I is pure 2-dimensional. Thus the map is injective and gives rise to a quotient $M_I = \mathcal{O}_X/\theta^r I$, with numerical invariants $(0, rC_m, 1)$. We claim that this family is generically stable (and thus in particular a stable sheaf exists!). To see this, note that M_I actually makes sense for an arbitrary sheaf $I \in \mathrm{Hilb}^1(X_{w^{-1},w^r})$; in the case of the ideal sheaf of a point x, the fact that $\mathrm{Hom}(\mathcal{O}_x, \mathcal{O}_X) = \mathrm{Ext}^1(\mathcal{O}_x, \mathcal{O}_X) = 0$ implies that $\mathrm{Hom}(\theta^r I_x, \mathcal{O}_X) \cong \mathrm{Hom}(\theta^r \mathcal{O}_X, \mathcal{O}_X) = \langle T^r\rangle$. The corresponding quotient M_{I_x} is not, in fact, stable, but *does* satisfy $R\,\mathrm{Hom}(N, M_{I_x}) = 0$ for every sheaf disjoint from C.

The assumption on w means that M_I (in the disjoint from C case) cannot have a proper nontrivial subsheaf of Chern class proportional to C_m (since there are no legal Euler characteristics), and thus if M_I is unstable, then it is destabilized by a sheaf of the form $\mathcal{O}_\alpha(f)$ for α in some finite set of roots. For each such effective α, let f_α be the minimum legal degree such that a subsheaf $\mathcal{O}_\alpha(f_\alpha)$ would destabilize M_I. Then any unstable M_I has a nonzero homomorphism from $\bigoplus_\alpha \mathcal{O}_\alpha(f_\alpha)$. Since

$$R\,\mathrm{Hom}(\bigoplus_\alpha \mathcal{O}_\alpha(f_\alpha), M_{I_x}) = 0 \tag{11.125}$$

for all $x \in C$, we conclude that

$$R\,\mathrm{Hom}(\bigoplus_\alpha \mathcal{O}_\alpha(f_\alpha), M_I) = 0 \tag{11.126}$$

for generic $I \in \mathrm{Hilb}^1(X_{w^{-1},w^r})$ as required.

Note that for w non-torsion, essentially the same argument shows not only that M_I is generically stable, but in fact that M_I is generically irreducible.

12 Derived equivalences

In the proof of Theorem 11.62, we constructed a number of derived equivalences between our noncommutative surfaces and commutative rational surfaces. Another such equivalence is suggested by Proposition 11.54 in comparison with [47, Thm. 7.1]: the relative Pic^r of a rational elliptic surface with an r-tuple fiber rC is given an explicit description there which is naturally isomorphic to a moduli space of the form of Proposition 11.54. In fact, we expect there to be a closer relation:

the relative Pic^r is not a fine moduli space, and thus gives rise to a natural sheaf of central simple algebras on the stable locus; it is thus natural to expect that sheaf of central simple algebras to be an Azumaya algebra of the form $X_{\rho;q;C}$.

Together, these results suggest that we should have a more general family of derived equivalences between surfaces of the form $X_{z,w}$. One such equivalence is straightforward.

Lemma 12.1. *There is an equivalence* $\Phi_{\mathcal{O}_{e_8}(-1)} : \operatorname{coh} X_{z/q,q} \cong \operatorname{coh} X_{z,q}$ *given by* $M \mapsto \theta M(-e_8)$.

It is instructive to consider how this equivalence interacts with the semiorthogonal decomposition corresponding to the blowdown from X_8 to X_7. Any object $M \in D^b \operatorname{coh} X_8$ fits into a distinguished triangle

$$N \otimes_k \mathcal{O}_{e_8}(-1) \to M \to L\alpha_8^! M' \to \tag{12.1}$$

where $M' \in D^b \operatorname{coh} X_7$ and $N \in D^b \operatorname{coh} \operatorname{Spec}(k)$. Applying $M \mapsto \theta M(-e_8)$ and using the relation between α_8^* and $\alpha_8^!$ gives a distinguished triangle

$$N \otimes_k \mathcal{O}_{e_8}(-1) \to \theta M(-e_8) \to L\alpha_8^! \theta M' \to \tag{12.2}$$

In other words, the given derived equivalence acts trivially on one part of the semiorthogonal decomposition, and acts as a shift of the induced Serre functor on the other part.

In general, we could attempt to do something similar replacing $\mathcal{O}_{e_8}(-1)$ by any exceptional object E. Indeed, for any exceptional object E, [10] tells us that there is a corresponding semiorthogonal decomposition, and an induced Serre functor on the kernel of $\mathcal{K}_E$ of $R\operatorname{Hom}(E, _)$. In our case, the induced Serre functor has the form $\theta_{\mathcal{K}_E}[2]$ where there is a functorial distinguished triangle

$$\theta_{\mathcal{K}_E} M \to \theta M \to R\operatorname{Hom}_k(R\operatorname{Hom}(M, E), \theta E) \to \tag{12.3}$$

(As usual, we use the fact that our category, as a derived category, has a natural dg-enhancement, so we have functorial cones.) We want to construct a derived functor by taking the distinguished triangle

$$E \otimes_k R\operatorname{Hom}(E, M) \to M \to N \to \tag{12.4}$$

with $N \in \mathcal{K}_E$ to a distinguished triangle of the form

$$E \otimes_k R\operatorname{Hom}(E, M) \to M \to \theta_{\mathcal{K}_E} N \to \tag{12.5}$$

Of course, the complication here is that the latter distinguished triangle lies in a different derived category with its own semiorthogonal decomposition. We would like the new derived category to also be of the form $D^b \operatorname{coh} X'_m$, and thus need to have a natural isomorphism

$$R\operatorname{Hom}_{X'_m}(\theta_{\mathcal{K}_E} M, E) \cong R\operatorname{Hom}_{X_m}(M, E) \tag{12.6}$$

We thus need to understand these two functors. In general, if $M \in \mathcal{K}_E$, then $R\operatorname{Hom}(M, \theta E) = R\operatorname{Hom}(E, M) = 0$, and thus $R\operatorname{Hom}(M, E) \cong R\operatorname{Hom}_C(M|_C^{\mathbf{L}}, E_C^{\mathbf{L}})$. Thus to understand the functor $R\operatorname{Hom}(\theta_{\mathcal{K}_E} _, E)$, we just need to understand $(\theta_{\mathcal{K}_E} _)|_C^{\mathbf{L}}$. Here, we have a distinguished triangle

$$(\theta_{\mathcal{K}_E} M)|_C^{\mathbf{L}} \to \theta M|_C^{\mathbf{L}} \to R\operatorname{Hom}_k(R\operatorname{Hom}(\theta M|_C^{\mathbf{L}}, \theta E|_C^{\mathbf{L}}), \theta E|_C^{\mathbf{L}}) \to \tag{12.7}$$

This, of course, we can compute entirely inside $D^b \operatorname{coh} C$. In fact, up to an application of θ (which on C simply twists by a line bundle), this is just a twist functor [57] in the spherical object $\theta E|_C^{\mathbf{L}}$. As a result, we can compute $(\theta_{\mathcal{K}_E} M)|_C^{\mathbf{L}}$ as the image of $M|_C^{\mathbf{L}}$ under a derived autoequivalence of $D^b \operatorname{coh} C$. But then to understand $R\operatorname{Hom}(\theta_{\mathcal{K}_E} M, E)$ for general M, we just need to compute the

image of $E|_C^{\mathbf{L}}$ under the inverse autoequivalence. This is easily computed as the sheaf F on C in the distinguished triangle

$$R\operatorname{Hom}_C(E|_C^{\mathbf{L}}, \theta^{-1}E|_C^{\mathbf{L}}) \otimes_k E|_C^{\mathbf{L}} \to \theta^{-1}E|_C^{\mathbf{L}} \to F \to \tag{12.8}$$

using the standard description of the inverse of a spherical twist functor.

Now, $X_{\rho;q;C}$ is naturally a sort of derived blowup of $\mathcal{K}_E$ in the sheaf $E|_C^{\mathbf{L}}$, and thus the above construction gives an alternate description as a derived blowup of $\mathcal{K}_E$ in the sheaf F. Most of the time, this cannot be simplified further, but if F is a deformation of $E|_C^{\mathbf{L}}$, we can hope to recognize the new description as a different noncommutative rational surface. This is what happens in the case $E = \mathcal{O}_{e_m}(-1)$, as the sheaf F is then just the structure sheaf of qx_m, giving a derived equivalence $D^b \operatorname{coh} X_{\rho;q;C} \cong D^b \operatorname{coh} X_{\rho;q;C}(-e_m)$ (which of course in that case is an actually abelian equivalence).

When $m = 8$, it turns out that the case $E = \mathcal{O}_X$ also works, giving us the following.

Lemma 12.2. *There is an equivalence $\Phi_{\mathcal{O}_X} : D^b \operatorname{coh} X_{z,qz} \cong D^b \operatorname{coh} X_{z,q}$ which takes $\mathcal{O}_X$ to $\mathcal{O}_X$ and on $\mathcal{K}_{\mathcal{O}_X}$ acts as $\kappa_q^{-1}\kappa_{qz}\theta_{\mathcal{K}_{\mathcal{O}_X}}$.*

Proof. By [40], to specify an equivalence between dg-categories equipped with semiorthogonal decompositions, it suffices to give equivalences between the two subcategories that are functorial for maps between the two subcategories. In other words, we just need to show that we have a functorial isomorphism

$$R\operatorname{Hom}(M, \mathcal{O}_X) \cong R\operatorname{Hom}(\kappa_q^{-1}\kappa_{qz}\theta_{\mathcal{K}_{\mathcal{O}_X}} M, \mathcal{O}_X). \tag{12.9}$$

for $M \in \mathcal{K}_{\mathcal{O}_X} \cong \mathcal{N}_{\rho;qz;C}$. But this reduces to the above calculation on C. Note that for a generic ineffective line bundle $\mathcal{L}$ of degree 0 on C, we have

$$\kappa_q^{-1}\kappa_{qz}\theta_{\mathcal{K}_{\mathcal{O}_X}} \mathcal{L} \cong \mathcal{L} \tag{12.10}$$

and thus the inverse of the relevant derived equivalence preserves $\mathcal{O}_C$ as required. □

Remark. Note that since Serre functors are unique, we can compute the Serre functor on $\mathcal{N}_{\rho;qz;C}$ by conjugating the induced Serre functor on $\mathcal{N}_{\rho;1;C}$ by κ_{qz}.

Combining the above operations gives the following.

Theorem 12.3. *For any element $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z})$, there is a derived equivalence $D^b \operatorname{coh} X_{z,q} \cong D^b \operatorname{coh} X_{z^a q^b, z^c q^d}$, which acts as an autoequivalence on $D^b \operatorname{coh} C$.*

Remark 1. Here we should caution the reader that this does not correspond to an action of $\mathrm{SL}_2(\mathbb{Z})$, as the relations of $\mathrm{SL}_2(\mathbb{Z})$ correspond to nontrivial autoequivalences (acting on the Grothendieck group in the same way as suitable powers of θ). If we allow contravariant equivalences, we can of course include the duality $R\operatorname{ad}$, extending $\mathrm{SL}_2(\mathbb{Z})$ to $\mathrm{GL}_2(\mathbb{Z})$. Note that since $R\operatorname{ad}$ is contravariant, it inverts the Serre functors, and thus each of $\Phi_{\mathcal{O}_X}$ or $\Phi_{\mathcal{O}_{e_8}(-1)}$ conjugates to its inverse under $R\operatorname{ad}$. We should also note that the abelian equivalence $X_{z,q} \cong X_{1/z,1/q}$ does *not* correspond to an element of the above family, since the corresponding autoequivalence of $D^b \operatorname{coh} C$ is the action of a hyperelliptic involution.

Remark 2. It is useful to record how the above operations act on $K_0(X_{z,q})$. Note first that for any root α of $E_8 \subset E_9$, the two generating equivalences preserve the class with numerical invariants $(0, \alpha, 0)$, as this class is in both $\mathcal{K}_{\mathcal{O}_{e_8}(-1)}$ and $\mathcal{K}_{\mathcal{O}_X}$ and is invariant under θ. We thus need

only determine how they act on the 4-dimensional orthogonal complement in $K_0(X_{z,q})$ of that 8-dimensional space, or in other words on classes of the form $(r, se_8 + tC_8, u)$. We find that the action on such classes of the derived equivalence $D^b \operatorname{coh} X_{z,q} \to D^b \operatorname{coh} X_{qz,q}$ is

$$\begin{aligned}(1,0,0) &\mapsto (1, -e_8 - C_8, 0)\\ (0,e_8,0) &\mapsto (0, e_8, 0)\\ (0,C_8,0) &\mapsto (0, C_8, -1)\\ (0,0,1) &\mapsto (0,0,1)\end{aligned}$$

while the action of $D^b \operatorname{coh} X_{z,q} \to D^b \operatorname{coh} X_{z,q/z}$ is

$$\begin{aligned}(1,0,0) &\mapsto (1, -C_8, 0)\\ (0,e_8,0) &\mapsto (1, e_8 - C_8, 0)\\ (0,C_8,0) &\mapsto (0, C_8, 0)\\ (0,0,1) &\mapsto (0, C_8, 1).\end{aligned}$$

The images of $(0, C_8, 0)$ and $(0,0,1)$ in each case follow immediately from the action on parameters, and the remaining degree of freedom (given that we must respect $\chi(_,_)$) is determined by the known fixed element $(0, e_8, 0)$, resp. $(1,0,1)$. For the general equivalence $D^b \operatorname{coh} X_{z,q} \to D^b \operatorname{coh} X_{z^a q^b, z^c q^d}$, we cannot be quite as precise (due to the aforementioned central extension); all we can say is that

$$(1,0,0) \mapsto (d, -be_8 + \frac{dh - b + c}{2} C_8, \frac{-bh - b - a + d}{2}) \tag{12.11}$$

$$(0,e_8,0) \mapsto (-c, ae_8 + \frac{-ch + c + a - d}{2} C_8, \frac{ah + b - c}{2}) \tag{12.12}$$

$$(0,C_8,0) \mapsto (0, dC_8, -b) \tag{12.13}$$

$$(0,0,1) \mapsto (0, -cC_8, a) \tag{12.14}$$

for some $h \in ab + bc + cd + 2\mathbb{Z}$ (corresponding to the application of a power of θ). Note that this gives a non-split extension of $\mathrm{SL}_2(\mathbb{Z})$ by $\mathbb{Z}$, split (by taking $h = 0$) over the index 2 (congruence) subgroup on which $ab + bc + cd \in 2\mathbb{Z}$.

One application of these equivalences is to finish up the proof of Theorem 11.62.

Corollary 12.4. *For any pair $r > 0$, d with $\gcd(r,d) = 1$, there exists a stable sheaf on X_{w^{-d}, w^r} with numerical invariants $(0, rC_m, d)$ and disjoint from C.*

Proof. Let $\Phi : D^b \operatorname{coh} X_{w,1} \cong D^b \operatorname{coh} X_{w^{-d}, w^r}$ be a derived equivalence of the form guaranteed by the Theorem. We claim that if x is a generic point of $X_{w,1}$, then (up to some constant shift) $\Phi \mathcal{O}_x$ is a stable sheaf with the desired numerical invariants and disjoint from C. Each of these conditions is an intersection of open conditions, so it will suffice to find points x_0, x_1 such that $\Phi \mathcal{O}_{x_0}$ is a stable sheaf with the correct invariants and $\Phi \mathcal{O}_{x_1}$ is disjoint from C. The latter is easy, since any derived equivalence between categories of our form respects restriction to C (up to an autoequivalence of $D^b \operatorname{coh} C$), so $\Phi \mathcal{O}_{x_1} |_C^{\mathbf{L}} = 0$ whenever $x_1 \notin C$. For the former, we take $x_0 \in C$. Since Φ acts as a derived autoequivalence on C, it follows that $\Phi \mathcal{O}_{x_0}$ is an even shift of a simple sheaf on C with rank r and degree d. Thus (fixing the shift as necessary) we conclude that $\Phi \mathcal{O}_{x_0}$ is a stable vector bundle on C, and thus stable as a sheaf on X_{w^{-d}, w^r}. □

Remark 1. Of course, the intersection of open sets where the above construction works will depend on the choice of stability condition. In particular, unlike the derived equivalence constructed in the proof of Theorem 11.62, we cannot expect that *every* point will map to a stable sheaf, especially in situations with multiple inequivalent stability conditions.

Remark 2. The construction given after the proof of Theorem 11.62 is actually a special case of the above construction; it is straightforward to apply the equivalence $D^b \coh X_{z,q} \to D^b \coh X_{z,z^{-r}q}$ to a generic point sheaf in the case $q = 1$, as long as $r < |\langle z \rangle|$. What happens when $r \geq |\langle z \rangle|$ is that the internal Serre functor begins to diverge from the usual Serre functor; if one replaces $\theta^l I$ by $\theta^l_{\mathcal{K}_{\mathcal{O}_X}} I$, the construction works in general.

We single out one particular composition of the above functors:

$$\Psi := \theta^{-1} \Phi_{\mathcal{O}_{e_8}(-1)} \Phi_{\mathcal{O}_X} \Phi_{\mathcal{O}_{e_8}(-1)} \cong (\Phi_{\mathcal{O}_X} \theta _(-e_8))(-e_8) \tag{12.15}$$

taking $D^b \coh X_{z,q} \to D^b \coh X_{q,1/z}$.

Lemma 12.5. *We have* $\Psi(\mathcal{O}_{e_8}) \cong \mathcal{O}_X$ *and* $\Psi(\mathcal{O}_X) \cong \mathcal{O}_{e_8}(-1)[-1]$.

Proof. Using the fact that θ commutes with any derived equivalence, we may rewrite the second claim using

$$\Psi = \Phi_{\mathcal{O}_{e_8}(-1)} \Phi_{\mathcal{O}_X} \theta^{-1} \Phi_{\mathcal{O}_{e_8}(-1)}, \tag{12.16}$$

and thus reduce to showing

$$\Phi_{\mathcal{O}_X}(\mathcal{O}_{e_8}) \cong \mathcal{O}_X(e_8), \tag{12.17}$$

$$\Phi_{\mathcal{O}_X}(\mathcal{O}_X(-e_8)) \cong \mathcal{O}_{e_8}(-1)[-1]. \tag{12.18}$$

These are equivalent via the distinguished triangle

$$\Phi_{\mathcal{O}_X}(\mathcal{O}_X(-e_8)) \to \Phi_{\mathcal{O}_X}(\mathcal{O}_X) \to \Phi_{\mathcal{O}_X}(\mathcal{O}_{e_8}) \to \tag{12.19}$$

with $\Phi_{\mathcal{O}_X}(\mathcal{O}_X) \cong \mathcal{O}_X$, so we need merely compute $\Phi_{\mathcal{O}_X}(\mathcal{O}_X(-e_8))$. For that, we note the distinguished triangle

$$\theta_{\mathcal{K}_{\mathcal{O}_X}} \mathcal{O}_X(-e_8) \to \theta \mathcal{O}_X(-e_8) \to \theta \mathcal{O}_X \to \tag{12.20}$$

so that $\mathcal{K}_{\mathcal{O}_X} \mathcal{O}_X(-e_8) \cong \mathcal{O}_{e_8}(-1)[-1]$; the action of $\kappa_q^{-1} \kappa_{qz}$ is then trivial to compute. □

This lets us prove a relation between our derived equivalences.

Corollary 12.6. *One has a natural isomorphism*

$$\Phi_{\mathcal{O}_{e_8}(-1)} \Phi_{\mathcal{O}_X} \Phi_{\mathcal{O}_{e_8}(-1)} \cong \Phi_{\mathcal{O}_X} \Phi_{\mathcal{O}_{e_8}(-1)} \Phi_{\mathcal{O}_X}. \tag{12.21}$$

Proof. Consider the two compositions

$$\Psi \circ \Phi_{\mathcal{O}_X} \qquad \text{and} \qquad \Phi_{\mathcal{O}_{e_8}(-1)} \circ \Psi. \tag{12.22}$$

Since $\Psi(\mathcal{O}_X) = \mathcal{O}_{e_8}(-1)[-1]$, we find that Ψ takes $\ker R\Hom(\mathcal{O}_X, _)$ to $\ker R\Hom(\mathcal{O}_{e_8}(-1), _)$. Since each of $\Phi_{\mathcal{O}_X}$ and $\Phi_{\mathcal{O}_{e_8}(-1)}$ acts as the internal Serre functor on the appropriate subcategory, and Serre functors are canonical, we conclude that the two compositions are naturally equivalent. Applying $\theta \Phi^{-1}_{\mathcal{O}_{e_8}(-1)}$ to both sides and simplifying gives the desired result. □

Remark. The action on $K_0(X_{z,w})$ suggests that $\theta^{-2}\Psi^4[2]$ should be the identity. If $2C_8 + e_8$ is ample, then this is true: $\theta^{-2}\Psi^4[2]$ preserves $\mathcal{O}_X$ and commutes with $\theta^{-1}\Phi_{\mathcal{O}_{e_8}(-1)}$ and thus preserves $\mathcal{O}_X(r(2C_8+e_8))$ for all r. Since it preserves a set of generators, it must be trivial. On the other hand, it is easy to see that if M is any nonzero sheaf with $M \cong \theta M$, $R\Gamma(M) = R\operatorname{Hom}(\mathcal{O}_{e_8}(-1), M) = 0$, then $\Psi(M) = M$, so $\theta^{-2}\Psi^4 M[2] \cong M[2] \ncong M$. Such a sheaf exists whenever there is a root α of E_8 with $\rho(\alpha) = 1$. Presumably one finds in such cases that $\theta^{-2}\Psi^4[2]$ is a composition of reflection functors, compare Theorem 12.20.

Another interesting consequence is that we can find derived categories of noncommutative surfaces with $K^2 = 1$ inside derived categories of *commutative* surfaces.

Corollary 12.7. *There is an isomorphism*

$$D^b \operatorname{coh} X_{\eta,x_0,\dots,x_7;q;C} \cong \mathcal{N}_{\eta,x_0,\dots,x_7,\eta^2 x_0^3/x_1\cdots x_7 q;1;C}. \tag{12.23}$$

Proof. We can think of $D^b \operatorname{coh} X_{\eta,x_0,\dots,x_7;q;C}$ via α_8^* as the subcategory of $D^b \operatorname{coh} X_{z,q}$ in which $R\operatorname{Hom}(\mathcal{O}_{e_8}, M) = 0$. Thus to establish the above isomorphism, we simply need to apply a derived equivalence of the form

$$D^b \operatorname{coh} X_{z,q} \to D^b \operatorname{coh} X_{z',q'} \tag{12.24}$$

taking $\mathcal{O}_{e_8}$ to $\mathcal{O}_X$. As we have seen, Ψ fills the bill. □

Remark 1. In fact, any $D^b \operatorname{coh} X_m$ is a triangulated subcategory of some $D^b \operatorname{coh} Y$ for a commutative variety Y, as follows from results of [40], but the resulting varieties Y have higher dimension than coming from the above construction, which induces an embedding in the derived category of a commutative surface whenever $m \le 7$.

Another case of interest starts with $X_{z,q}$ with q torsion of order l. In that case, we can ask how the derived equivalences act on simple 0-dimensional sheaves of degree l. The argument of Corollary 12.4 tells us that the result will generically be a semistable sheaf disjoint from C, and gives us the following consequence.

Proposition 12.8. *Suppose $ad - rb = 1$, with $r > 0$. Then for q of order l and arbitrary z, the moduli space of stable sheaves on $X_{z^d q^{-b}, z^{-r} q^a}$ with invariants $(0, lrC_8, ld)$ is nonempty. Moreover, the obstruction to the existence of a universal family over the generic point of the resulting component of the moduli space is given by the generic fiber of the Azumaya algebra $X_{z,q}$.*

Proof. The one thing we need to argue is that the sheaves are actually generically stable, not just semistable. The only way a sheaf disjoint from C with these numerical invariants can fail to be irreducible is if it has a subsheaf with Chern class a root of $W(E_9)$, and only finitely many such roots can appear. Since the sheaves supported on C do not have morphisms to or from such root sheaves, the generic sheaf also has no such morphisms. □

This gives us a generalization of [47, Thm. 7.1] (or, more precisely, the special case of that result with the additional hypothesis that the anticanonical curve is smooth; we will remove that hypothesis in [44]). The idea here is that if z is l-torsion, then the linear system $|lC_8|$ on $X_{z,1}$ is an elliptic pencil with a fiber of the form lC. Then for any degree d, we may compute the relative Pic^d of this fibration, and the result will have a canonical minimal proper regular model, which we can hope to identify.

Corollary 12.9. *Let z be an l-torsion point of $\mathrm{Pic}^0(C)$, let d be an integer, and let $\phi : C \to C'$ be the $\gcd(l,d)$-isogeny with kernel $\langle z^{l/\gcd(l,d)} \rangle$. Then the minimal proper regular model of the relative Pic^d of the elliptic surface*

$$X_{\eta,x_0,\ldots,x_7,\eta^2 x_0^3/x_1\cdots x_7 z;1;C} \tag{12.25}$$

is isomorphic to

$$X_{\phi(\eta),\phi(x_0),\ldots,\phi(x_7),\phi(\eta^2 x_0^3/x_1\cdots x_7 z^a);1;C'}, \tag{12.26}$$

where a is such that $ad \cong \gcd(l,d) \bmod l$.

Proof. Set $X := X_{z,1}$, $Y := X_{z^a, z^{l/\gcd(l,d)}}$, and let Z denote the putative minimal proper regular model, a.k.a. the center of Y. The constraint on a implies the existence of a unique matrix

$$\begin{pmatrix} a & b \\ l/\gcd(l,d) & d/\gcd(l,d) \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}) \tag{12.27}$$

and thus a corresponding derived equivalence $D^b \operatorname{coh} Y \cong D^b \operatorname{coh} X$.

Now let M_X be the moduli space of simple sheaves on X with invariants $(0, lC_8, d)$, and let M_Y be the moduli space of simple sheaves on Y with invariants $(0, 0, \gcd(l,d))$. The derived equivalence has the correct effect on invariants to map sheaves in M_Y to sheaves in M_X, and the usual argument tells us that it generically takes sheaves to sheaves (up to an overall shift which we can correct if necessary) and respects disjointness from C. In fact something stronger is true: since Serre functors of derived categories are unique, the derived equivalence commutes with θ and thus with θ^l, so in particular identifies the spaces of natural transformations $\theta^l \to \mathrm{id}$ on either side. On the commutative surface X, such natural transformations are in natural bijection with the global sections of ω_X^{-l}, and thus, mod scalars, correspond to fibers of the elliptic fibration on X. For Y, we note that the global sections of the anticanonical bundle on Z induce such natural transformations, and thus by dimensionality exhaust the space of such global sections. In particular, we have a natural identification between the fibers of X and the fibers of Z. We thus find that the image of a sheaf supported on a given fiber of Z maps via the derived equivalence to an object supported on the corresponding fiber of X. (That is, the corresponding morphism $\theta^l M \to M$ is 0.)

In particular, given any point of Z not on C, we may take a simple module over the corresponding fiber of the Azumaya algebra and apply the derived equivalence, and the result will generically be a line bundle of degree d on the corresponding fiber of X. As a result, the derived equivalence identifies the generic point of the generic fiber of Z with the generic point of Pic^d of the generic fiber of X. Since a smooth curve is determined by its function field, this is enough to determine the entirety of Pic^d of the generic fiber of X. Since Z is a minimal proper regular elliptic fibration with the same generic fiber as the relative Pic^d of X, it must in fact be the minimal proper regular model as required. □

Remark. We could, of course, also consider the relative moduli space of stable vector bundles of rank r and degree d, but this is generically the same as the relative Pic^d, so has the same minimal proper regular model. We should also note here that $\phi(z^a)$ is independent of the choice of a, since by assumption $\phi(z^a)^{d/\gcd(l,d)} = \phi(z)$ and $d/\gcd(l,d)$ is relatively prime to $|\ker \phi|$.

Since we have a reasonable understanding of the Hilbert scheme, it is natural to investigate how the Hilbert scheme interacts with the derived equivalences, in particular how the equivalences act on "ideal sheaves". Most of the time, this will give sheaves of high rank, but there are a couple of cases where things are tractable. The first of these is when the image again has rank 1, giving maps between Hilbert schemes. If $I \in \mathrm{Hilb}^n(X_{z,q})$, then the derived equivalence $\Phi : D^b \operatorname{coh} X_{z,q} \to$

$D^b \operatorname{coh} X_{z,q/z}$ produces an object $\Phi(I)$ with numerical invariants $(1, -(n-1)C_8, 1-n)$, and thus $\theta^{1-n}\Phi(I)$ has the same numerical invariants as an object in $\operatorname{Hilb}^n(X_{z,q/z})$.

Proposition 12.10. *If $n < \max(|\langle q/z\rangle|, |\langle q\rangle|)$, then the above construction induces a birational map $\operatorname{Hilb}^n(X_{z,q}) \dashrightarrow \operatorname{Hilb}^n(X_{z,q/z})$.*

Proof. Suppose I is given as the kernel of a surjection $I' \to \mathcal{O}_y$ with $y \in C$ and I' a generic point of $\operatorname{Hilb}^{n-1}(X_{z,q})$. Since $\Phi(I')$ is a point in $\operatorname{Hilb}^{n-1}(X_{z,q/z})$ by induction, and $\Phi(\mathcal{O}_y)$ is a line bundle of degree 1, we find that $\Phi(I)$ is represented by a complex

$$\Phi(I') \to \mathcal{L}. \tag{12.28}$$

Taking the adjoint represents $R\operatorname{ad}\Phi(I)$ as an extension

$$0 \to \operatorname{ad}\Phi(I') \to R\operatorname{ad}\Phi(I) \to R^1 \operatorname{ad}\mathcal{L} \to 0, \tag{12.29}$$

where $R\operatorname{ad}\Phi(I')$ is a sheaf as long as $n \le |\langle q/z\rangle|$, per Proposition 11.53. In particular, $R\operatorname{ad}\Phi(I)$ is a sheaf, which is pure 2-dimensional since the map $\Phi(I') \to \mathcal{L}$ is nonzero, so has 0-dimensional cokernel.

It follows that for generic I, $R\operatorname{ad}\Phi(I)$ is a pure 2-dimensional sheaf, and thus, again using Proposition 11.53, the same applies to $\Phi(I)$.

The analogous calculation with $\Phi R\operatorname{ad}$ instead of $R\operatorname{ad}\Phi$ (applied to a sheaf in $\operatorname{Hilb}^n(X_{\rho;1/q;C})$) gives a proof valid when $n < |\langle q\rangle|$. □

Remark. This presumably corresponds to the usual $c \mapsto c+1$ symmetry of Calogero-Moser spaces.

The other case of interest is when the image has rank 0. We find that if $I \in \operatorname{Hilb}^n(X_{z,q})$, then ΨI has numerical invariants $(0, -nC_8 - e_8, 0)$, and thus we expect generically to obtain a shift of a semistable sheaf with invariants $(0, nC_8 + e_8, 0)$. To guarantee the sheaf is semistable, it is simplest to ensure that it is in the kernel of $R\Gamma$, or equivalently that I is in the image of α_8^*.

To compute ΨI in this case, we note that $\theta^n I(-e_8)$ has numerical invariants $(1, -nC_8 - e_8, 0)$. Assuming this is in $\mathcal{K}_{\mathcal{O}_X}$ (which we know happens generically), applying $\Phi_{\mathcal{O}_X}$ reduces to the internal Serre functor, which in turn generically gives the cone of the morphism

$$\theta^{n+1} I(-e_8) \to \theta\mathcal{O}_X. \tag{12.30}$$

We thus obtain a short exact sequence

$$0 \to \theta^2 I(-2e_8) \to \theta^{2-n}\mathcal{O}_X(-e_8) \to \Psi I[1] \to 0. \tag{12.31}$$

which we thus see is closely related to the relation of the previous section between the Hilbert scheme and 1-dimensional sheaves.

This is a noncommutative elliptic analogue of the construction of [18, Defn. 5.2.3], which constructed a generalized Calogero-Moser space from certain Deligne-Simpson problems which by the results of [47] can be translated into moduli spaces of sheaves on rational surfaces; up to admissible reflections, these are degenerations of the above moduli problem with $z = 1$. In the case $q = z = 1$, the moduli space can be explicitly computed [22], and is in general (when there are effective roots) a perverse Hilbert scheme rather than a Hilbert scheme; it is unclear therefore whether the above birational map is regular for generic q, z. Of course, to understand this would require better control over the action of the derived equivalences on the t-structure.

The use of derived equivalences in proving Theorem 11.62 above relied crucially on the fact that we could recover the surface from the restriction of K_0 to C. This turns out to be an even stronger constraint for $m \neq 8$, and we have the following.

Lemma 12.11. *Suppose $\Phi : D^b \operatorname{coh} X_{\rho;q;C} \cong D^b \operatorname{coh} X_{\rho';q';C'}$ is a derived equivalence. Then $m = m'$, $C \cong C'$, and the derived equivalence restricts to a derived equivalence $D^b \operatorname{coh} C \cong D^b \operatorname{coh} C'$. If $m \neq 8$, then this derived equivalence preserves rank and the isomorphism $C \cong C'$ can be chosen to identify q and q'.*

Proof. That we obtain a derived equivalence $D^b \operatorname{coh} C \cong D^b \operatorname{coh} C'$, and thus $C \cong C'$, follows as in the proof of Theorem 11.62; that $m = m'$ follows from the fact that both surfaces have isomorphic K_0 along with the fact that $\operatorname{rank}(K_0(X)) = m + 4$. We furthermore have $\Phi(M)|^{\mathbf{L}}_{C'} \cong \Phi(M|^{\mathbf{L}}_C)$ for any object M. In particular, we observe that both $[\mathcal{O}_C]|_C$ and $[\mathcal{O}_{C'}]|_{C'}$ have rank 0 and degree $m-8$, so that when $m \neq 8$, Φ must preserve the class of a point, and thus respects $q = [pt]|_C \in K_0(C)$. Since it also respects the χ pairing, it preserves rank. □

Note that for $m = 8$, if a derived equivalence Φ does not preserve the class of a point, then for any x, $\Phi\mathcal{O}_x$ will be a shift of a vector bundle of rank r and degree d with $\gcd(r, d) = 1$. But then there is a derived equivalence coming from Theorem 12.3 taking $\Phi\mathcal{O}_x$ back to a shift of a point. Thus for any m the problem of understanding derived equivalences reduces to that of understanding derived equivalences taking points to points, and thus (applying an element of $\operatorname{Aut}(C)$ as appropriate) such that $\Phi\mathcal{O}_x \cong \mathcal{O}_x$ for all $x \in C$. (Then Φ acts on $D^b \operatorname{coh} C$ as the twist by a line bundle.)

Since our categories have exceptional collections, it will suffice to understand how the exceptional objects in those collections behave under such derived equivalences. The assumption on Φ ensures that if $\operatorname{rank}(E) = 1$, then $\Phi E|^{\mathbf{L}}_C$ is a line bundle, while if $\operatorname{rank}(E) = 0$, then $\Phi E|^{\mathbf{L}}_C$ is the structure sheaf of a point. Since $|^{\mathbf{L}}_C$ has homological dimension 1, we conclude that $h^0(h^p\Phi E)|^{\mathbf{L}}_C = 0$ unless $p = 0$ and $h^{-1}(h^p\Phi E)|^{\mathbf{L}}_C = 0$ unless $p = 1$. This implies that $h^0\Phi E$ is transverse to C and $h^p\Phi E$ is disjoint from C for $p \neq 0$.

To understand the structure further, we have the following result.

Lemma 12.12. *Let E be an exceptional object of $D^b \operatorname{coh} X_{\rho;q;C}$. Then $\operatorname{Ext}^1(h^p E, h^q E) = 0$ for all p, q.*

Proof. In [66, III(4.6.10.1)], Verdier constructed a spectral sequence

$$E_2^{p,q} = \prod_{d\in\mathbb{Z}} \operatorname{Ext}^p(h^{d-q}M, h^d N) \Rightarrow \operatorname{Ext}^{p+q}(M, N) \tag{12.32}$$

in the derived category of an abelian category with enough injectives, where the cohomology of M is bounded above and N can be represented by a complex which is bounded below. In our case, we can represent both M and N by bounded complexes, and we may compute inside $D \operatorname{qcoh} X_{\rho;q;C}$, so there is no difficulty applying this for $M, N \in D^b \operatorname{coh} X_{\rho;q;C}$. In particular, for $M = N = E$, we immediately conclude (since the sequence collapses on the E_3 page, and there are no nonzero maps hitting $E_2^{1,q}$) that

$$\prod_{d\in\mathbb{Z}} \operatorname{Ext}^1(h^{d-q}E, h^d E) = 0 \tag{12.33}$$

for $q \neq -1$. For $q = -1$, we simply note that $k = \operatorname{End}(E)$ contributes to the $E_\infty^{0,0}$ term of the limit, and thus $E_2^{1,-1} = E_\infty^{1,-1}$ must still vanish. □

Remark. This was shown in [27] for spherical objects, but the proof is essentially the same.

Lemma 12.13. *Suppose we have an exact sequence $0 \to T \to M \to F \to 0$ such that $\operatorname{Hom}(T, F) = \operatorname{Ext}^2(F, T) = \operatorname{Ext}^1(M, M) = 0$. Then $\operatorname{Ext}^1(T, T) = \operatorname{Ext}^1(F, F) = 0$.*

Proof. The composition

$$\mathrm{Ext}^1(M,T) \to \mathrm{Ext}^1(T,T) \to \mathrm{Ext}^1(T,M) \tag{12.34}$$

factors through $\mathrm{Ext}^1(M,M) = 0$ and thus vanishes. On the other hand, the first map has cokernel $\mathrm{Ext}^2(F,T) = 0$, so is surjective and the second map has kernel $\mathrm{Hom}(T,F) = 0$, so is injective; it follows that $\mathrm{Ext}^1(T,T) = 0$.

The dual argument gives $\mathrm{Ext}^1(F,F) = 0$. □

Let $S_{\rho;q;C}$ be the set of -2-curves of $X_{\rho;q;C}$; i.e., $\alpha \in \mathrm{Pic}(X)$ such that there is an irreducible 1-dimensional sheaf disjoint from C with Chern class α. As in the commutative case, this is the set of simple roots of a root system $\Phi(S_{\rho;q;C})$ defined as the set of real roots α of E_{m+1} such that $\rho(\alpha) \in q^{\mathbb{Z}}$. Call a sheaf of the form $\mathcal{O}_\alpha(d)$ with α a -2-curve a "-2-sheaf", and similarly for "-1-sheaf".

Lemma 12.14. *Suppose $M \in \mathrm{coh}\, X_{\rho;q;C}$ satisfies $M|_C = 0$, $\mathrm{Ext}^1(M,M) = 0$. Then M admits a filtration in which every subquotient is a -2-sheaf.*

Proof. Note that M cannot be 0-dimensional, since 0-dimensional sheaves always admit deformations. More generally, if T is the maximal 0-dimensional subsheaf of M, then Lemma 12.13 implies $\mathrm{Ext}^1(T,T) = 0$ and thus $T = 0$, so that M is pure 1-dimensional. In addition, $c_1(M) \cdot e_m \geq 0$, since otherwise it would have a sub- or quotient sheaf of the form $\mathcal{O}_{e_m}(d)$, contradicting disjointness from C.

Now, if M is any pure 1-dimensional sheaf with $c_1(M) = l\alpha$ for some $\alpha \in S_{\rho;q;C}$ (ignoring the rigidity condition $\mathrm{Ext}^1(M,M) = 0$), then $\alpha \cdot c_1(M) < 0$, and thus there is a morphism between M and $\mathcal{O}_\alpha(d)$ for any d such that the latter exists. A map $M \to \mathcal{O}_\alpha(d)$ has image $\mathcal{O}_\alpha(d')$ and kernel which by induction has a filtration as required. If q is torsion, then we can always increase d until there is no map from $\mathcal{O}_\alpha(d)$ and thus end up in this case. Otherwise q is non-torsion, and then disjointness from C implies that the cokernel of a map from $\mathcal{O}_\alpha(d)$ is pure 1-dimensional, so that we may again induct.

More generally, the condition $\mathrm{Ext}^1(M,M) = 0$ implies $\chi(M,M) > 0$ and thus $c_1(M)^2 < 0$. Since $c_1(M)$ is effective, it cannot be nef, and thus (since it satisfies $c_1(M) \cdot C_m = 0$, $c_1(M) \cdot e_m \geq 0$) there is some simple root σ having negative intersection with $c_1(M)$. If that root is ineffective, we may apply the corresponding reflection. Since the process cannot otherwise terminate, we will reach a point where the simple root is effective, and the corresponding divisor class on the original surface will be in $S_{\rho;q;C}$. In other words, there is some $\alpha \in S_{\rho;q;C}$ such that $\alpha \cdot c_1(M) < 0$. Now, let T be the maximal subsheaf of M which has Chern class a multiple of α. Then T is disjoint from C with $\mathrm{Hom}(T, M/T) = 0$, so that Lemma 12.13 gives $\mathrm{Ext}^1(M/T, M/T) = 0$, allowing us to induct. □

Lemma 12.15. *Let M be a coherent sheaf on $X_{\rho;q;C}$ such that $\mathrm{Ext}^1(M,M) = 0$ and $M|_C^{\mathbf{L}}$ is a line bundle. Then the torsion subsheaf of M is an extension of -2-sheaves and the torsion-free quotient is a line bundle.*

Proof. If $M|_C$ is a line bundle, so $\mathrm{rank}(M) = 1$, let T be the torsion subsheaf of M. Since M/T is pure 2-dimensional, it is transverse to C, and thus we have a short exact sequence

$$0 \to T|_C \to E|_C \to (M/T)|_C \to 0. \tag{12.35}$$

Since $M|_C$ is a line bundle and $\mathrm{rank}(T|_C) = 0$, we conclude that T is disjoint from C. But then Lemma 12.13 implies $\mathrm{Ext}^1(T,T) = \mathrm{Ext}^1(M/T, M/T) = 0$. The claim about the torsion subsheaf then follows from Lemma 12.14, while the claim for M/T follows from Corollary 11.36, since $\mathrm{rank}(M/T) = \mathrm{rank}(M) = \mathrm{rank}(M|_C^{\mathbf{L}}) = 1$. □

Lemma 12.16. *Let M be a coherent sheaf on $X_{\rho;q;C}$ such that $\mathrm{Ext}^1(M,M) = 0$ and $M|_C^{\mathbf{L}}$ is the structure sheaf of a point. Then M has a filtration in which all but one subquotient is a -2-sheaf and the remaining subquotient is a -1-sheaf.*

Proof. Lemma 12.13 allows us to peel off the maximal subsheaf which is an extension of -2-sheaves, and then the maximal quotient which is such an extension, reducing to the case that M has no maps to or from -2-sheaves. Now, $\mathrm{rank}(M) = \mathrm{rank}(M|_C^{\mathbf{L}}) = 0$, so $c_1(M) < 0$, and thus $c_1(M)$ is not nef. Since $c_1(M)\cdot\alpha < 0$ for all -2-curves α and $c_1(M)\cdot C_m > 0$, the only way the algorithm for testing nefness of $c_1(M)$ can terminate is with $wc_1(M)\cdot e_m < 0$. This implies that there is a map between M and any sheaf of the form $\mathcal{O}_e(d)$, where e is the corresponding -1-curve. In particular, M has an irreducible constituent of this form, and the residual sheaf is rigid and disjoint from C, so (since M has no maps to or from -2-sheaves) must vanish. □

If q is not torsion, then each $\alpha \in S_{\rho;q;C}$ corresponds to a unique sheaf $\mathcal{O}_\alpha(d_\alpha)$, and we may then define a *reflection functor* R_α on $D^b \operatorname{coh} X_{\rho;q;C}$ as the cone

$$\mathcal{O}_\alpha(d_\alpha) \otimes_k R\operatorname{Hom}(\mathcal{O}_\alpha(d_\alpha), M) \to M \to R_\alpha M \to, \tag{12.36}$$

that is, the twist functor in the sense of [57] in the spherical object $\mathcal{O}_\alpha(d_\alpha)$. Note

$$R_\alpha \mathcal{O}_\alpha(d_\alpha) = \mathcal{O}_\alpha(d_\alpha)[-1], \tag{12.37}$$

and we have an inverse functor fitting into the distinguished triangle

$$R_\alpha^{-1} M \to M \to \mathrm{Hom}_k(R\operatorname{Hom}(M, \mathcal{O}_\alpha(d_\alpha)), \mathcal{O}_\alpha(d_\alpha)) \to . \tag{12.38}$$

The two constructions are dual, and indeed $R_\alpha R\operatorname{ad} = R\operatorname{ad} R_\alpha^{-1}$. In addition, if α and β are two distinct roots of $S_{\rho;q;C}$ then we have the braid relations [57]

$$\begin{aligned} R_\alpha R_\beta &\cong R_\beta R_\alpha & \alpha\cdot\beta &= 0 \\ R_\alpha R_\beta R_\alpha &\cong R_\beta R_\alpha R_\beta & \alpha\cdot\beta &= 1, \end{aligned} \tag{12.39}$$

so that the reflection functors and their inverses induce a representation of the associated braid group. In particular, for any element w of the corresponding Coxeter group, we can take the composition of reflection functors corresponding to any reduced word for w, and the result $R(w)$ will be independent of the choice of reduced word.

Lemma 12.17. *Suppose q is not torsion. Then for all $w \in W(S_{\rho;q;C})$, $R(w)$ preserves the nonnegative part of the t-structure, and $R(w)^{-1}$ preserves the nonpositive part of the t-structure. Moreover, if $M = \tau_{\geq 0} M$ is such that for any -2-sheaf N, $\mathrm{Hom}(N, M) = 0$, then the same holds for $R(w)M$. Dually, if $M = \tau_{\leq 0} M$ is such that for any -2-sheaf N, $\mathrm{Hom}(M, N) = 0$, then the same holds for $R(w)^{-1}M$.*

Proof. This reduces immediately to the case that w is a simple reflection, and via the adjoint symmetry to the case of R_α. If q is not torsion, then any subsheaf of $\mathcal{O}_\alpha(d_\alpha)^l$ has the same form, and thus the natural map

$$\mathcal{O}_\alpha(d_\alpha) \otimes_k \mathrm{Hom}(\mathcal{O}_\alpha(d_\alpha), M) \to h^0(M) \tag{12.40}$$

is necessarily injective, so that $h^p R_\alpha M = 0$ for $p < 0$. Moreover, if $\mathrm{Hom}(\mathcal{O}_\alpha(d_\alpha), M) = 0$, then $h^0(R_\alpha M)$ is an extension of $\mathcal{O}_\alpha(d_\alpha)$ by M, so still has no map from a -2-sheaf other than $\mathcal{O}_\alpha(d_\alpha)$, for which we have

$$\mathrm{Hom}(\mathcal{O}_\alpha(d_\alpha), R_\alpha M) \cong \mathrm{Hom}(R_\alpha^{-1}\mathcal{O}_\alpha(d_\alpha), M) \cong \mathrm{Hom}(\mathcal{O}_\alpha(d_\alpha)[1], M) = 0. \tag{12.41}$$

□

Remark. Note that R_α is actually the derived functor of $h^0 R_\alpha$, but this fails for $\ell(w) > 1$; if $\alpha \cdot \beta \neq 0$, then $h^0 R_\alpha$ does not take injective objects to R_β-acyclic objects in general, so the composition of the derived functors is not the derived functor of the composition.

Lemma 12.18. *Suppose q is non-torsion and the root system $\Phi(S_{\rho;q;C})$ is finite. Let M be an extension of -2-sheaves, and define a sequence $M_0 = M, M_1, \ldots, M_n$ of objects such that for each i we have $M_i = R_\alpha M_{i-1}$ for some α such that $\mathrm{Hom}(\mathcal{O}_\alpha(d_\alpha), M_{i-1}) \neq 0$. Then $n \leq |\Phi(S_{\rho;q;C})|$ and $h^0(M_n) = 0$.*

Proof. Let w_i be the corresponding sequence of Weyl group elements such that $M_i = R(w_i)M$, with $w_i = s_{\alpha_i} w_{i-1}$. It will suffice to show that $\ell(w_i) = \ell(w_{i-1}) + 1$, since then the process will be forced to terminate after at most $|\Phi(S_{\rho;q;C})|$ steps and the only way it can terminate is when $h^0(M_n) = 0$.

To see that the length must increase, simply note that for any expression $w_{i-1} = s_\alpha w'$, we have

$$\mathrm{Hom}(\mathcal{O}_\alpha(d_\alpha), R(w_{i-1})M) \cong \mathrm{Hom}(R_\alpha^{-1}\mathcal{O}_\alpha(d_\alpha), R(w')M) \cong \mathrm{Hom}(\mathcal{O}_\alpha(d_\alpha)[1], R(w')M) = 0, \tag{12.42}$$

and thus $\alpha_i \neq \alpha$. □

Remark. Note that there are instances with $|\Phi(S_{\rho;q;C})| = \infty$ where this fails even for sheaves with $\mathrm{Ext}^1(M, M) = 0$. Indeed, suppose α and β are two -2-curves with $l = \alpha \cdot \beta > 1$, and consider the sheaf $M = R_\beta \mathcal{O}_\alpha(d_\alpha)$. This has no maps from $\mathcal{O}_\beta(d_\beta)$, so our only choice is to apply R_α, and (since the length must increase) alternate between that and R_β from then on. One can show that M_n is always a sheaf, and thus invertibility prevents the process from terminating. (The argument reduces to showing that the generic quadratic algebra in l variables with a single relation is Koszul, with Hilbert series $1/(1 - lt + t^2)$.)

Lemma 12.19. *Suppose $S_{\rho;q;C}$ is empty or q is non-torsion and the root system $\Phi(S_{\rho;q;C})$ is finite. If $E \in D^b \operatorname{coh} X_{\rho;q;C}$ is an exceptional object such that $E|_C^{\mathbf{L}}$ is a the structure sheaf of a point, then there is an element Ψ of the group generated by reflection functors such that ΨE is a -1-sheaf.*

Proof. It suffices to construct Ψ such that ΨE is a sheaf with no maps to or from -2-sheaves, since then Lemma 12.16 implies the desired result.

In general, we know that $h^p E$ for $p \neq 0$ is disjoint from C and has no self-Ext^1, so is an extension of -2-sheaves. If $S_{\rho;q;C}$ is empty, then E is already a sheaf and we are done.

Otherwise, q is non-torsion and $S_{\rho;q;C}$ generates a finite root system. An easy induction using the adjoint of Lemma 12.18 tells us that for any object of $D^b \operatorname{coh} X_{\rho;q;C}$ such that every cohomology sheaf is an extension of -2-sheaves, there is a finite composition of inverse reflection functors taking the object to one supported only in negative degrees.

Now, E fits into a distinguished triangle

$$T \to E \to F \to \tag{12.43}$$

such that $\mathrm{Ext}^p(T, N) = 0$ for any -2-sheaf N and any $p \leq 0$. If Ψ_1 is a product of inverse reflection functors such that $\tau_{\geq 0}\Psi_1 F = 0$, then we have a distinguished triangle

$$\Psi_1 T \to \Psi_1 E \to \Psi_1 F \to \tag{12.44}$$

from which we may conclude that $\tau_{\geq 1}\Psi_1 E = 0$ and $\mathrm{Hom}(\Psi_1 E, N) = 0$ for all -2-sheaves N. We may thus reduce to the case that $E = \tau_{\leq 0} E$ and $h^0(E)$ has no maps to -2-sheaves.

Now, let p be minimal such that $h^p E \neq 0$. If there is no -2-sheaf N with $\mathrm{Hom}(N, h^p E) \neq 0$, then $p = 0$ and we are done. Otherwise, if we can show that $\mathrm{Ext}^1(N, h^0 E) = 0$ for any such sheaf,

then applying the corresponding reflection functor will not affect the $\mathrm{Ext}^p(E,N)=0$ condition and it will then follow easily from Lemma 12.18 that this process will terminate as required in an exceptional sheaf with no maps to or from -2-curves.

Consider the short exact sequence

$$0 \to N^l \to h^p E \to M \to 0, \tag{12.45}$$

with l maximal. Since $\mathrm{Ext}^1(h^pE, h^0E)=0$, the long exact sequence reads in part

$$0 \to \mathrm{Ext}^1(N^l, h^0E) \to \mathrm{Ext}^2(M, h^0E), \tag{12.46}$$

so that it suffices to prove vanishing of $\mathrm{Ext}^2(M,h^0E)$. If $p<0$, then $\theta M = M$ and thus $\mathrm{Ext}^2(M,h^0E)\cong \mathrm{Hom}(h^0E,M)^*$, which vanishes since M is an extension of -2-sheaves, to which h^0E has no maps.

Now, for $p=0$, Lemma 12.16 tells us that E is an extension of a -1-sheaf E' by an extension of -2-sheaves. In particular, this tells us that any subsheaf of E which is disjoint from C is an extension of -2-sheaves, and the same holds for M. In particular, we have a short exact sequence

$$0 \to M_1 \to M \to E' \to 0 \tag{12.47}$$

where M_1 is an extension of -2-sheaves.

Since $\mathrm{Ext}^2(M,N)\cong \mathrm{Hom}(N,M)^*=0$, we have $\mathrm{Ext}^2(M,E)\cong \mathrm{Ext}^2(M,M)$. But this vanishes, since we have the exact sequences

$$\mathrm{Ext}^2(E',M) \to \mathrm{Ext}^2(M,M) \to \mathrm{Ext}^2(M_1,M)=0 \tag{12.48}$$

and

$$0=\mathrm{Ext}^2(E',M_1) \to \mathrm{Ext}^2(E',M) \to \mathrm{Ext}^2(E',E')=0, \tag{12.49}$$

using $\mathrm{Ext}^2(M_1,M)\cong \mathrm{Hom}(M,M_1)^*=0$ and $\mathrm{Ext}^2(E',M_1)\cong \mathrm{Hom}(M_1,E')^*=0$. □

Theorem 12.20. *Suppose either $S_{\rho;q;C}$ is empty or $|\langle q\rangle|=\infty$ and $\Phi(S_{\rho;q;C})$ is finite. Then any derived equivalence*

$$\Psi : X_{\rho;q;C} \to X_{\rho';q;C} \tag{12.50}$$

such that $\Psi\mathcal{O}_x=\mathcal{O}_x$ for all $x\in C$ is in the group generated by twists by line bundles, admissible reflections, and reflection functors.

Proof. Suppose $m\geq 1$, and consider the exceptional object $\Psi\mathcal{O}_{e_m}(-1)$. Lemma 12.19 tells us that there is an element Ψ_1 of the group generated by reflection functors such that

$$\Psi_1\Psi\mathcal{O}_{e_m}(-1) \tag{12.51}$$

is a -1-sheaf. But then there is an admissible element $w\in W(E_{m_1})$ such that

$$w\Psi_1\Psi\mathcal{O}_{e_m}(-1)\cong \mathcal{O}_{e_m}(d) \tag{12.52}$$

for some d, at which point twisting by $(d+1)e_m$ gives $\mathcal{O}_{e_m}(-1)$. We thus reduce to proving the Theorem in the case $\Phi\mathcal{O}_{e_m}(-1)=\mathcal{O}_{e_m}(-1)$, which then immediately reduces to the corresponding claim for X_{m-1}.

This similarly works to reduce from X_0' to X_{-1}, replacing e_m by s, and thus leaves us only the X_{-1} and X_0 cases to consider.

For X_{-1}, there are no -2-curves, so that $\Psi\mathcal{O}_X$ is already a line bundle by Lemma 12.15, and we may as well assume $\Psi\mathcal{O}_X \cong \mathcal{O}_X$. But then $\Psi\theta^{-l}\mathcal{O}_X \cong \theta^{-l}\Psi\mathcal{O}_X \cong \theta^{-l}\mathcal{O}_X$ for all l, and thus Ψ actually preserves the t-structure. Now Euler characteristic considerations force $\Psi\mathcal{O}_X(-d) \cong \mathcal{O}_X(-d)$ for $d \in \{0,1,2\}$, and since restriction to C is faithful on that three-object subcategory (which generates $D^b \operatorname{coh} X$), this forces Ψ to be naturally isomorphic to the identity.

For X_0, let x_1 be any point of C, and observe that any derived equivalence between $X_{\eta,\eta';q;C}$ and $X_{\bar\eta,\bar\eta';q;C}$ extends to a derived equivalence on the blowup that fixes $\mathcal{O}_{e_1}(-1)$ (relative to the "even" basis of $\operatorname{Pic}(X)$). Since the stabilizer of $\mathcal{O}_{e_1}(-1)$ contains the reflection (whether admissible or a reflection functor), we easily verify that it is the correct group. □

Remark. Note that the finiteness hypothesis is automatically satisfied for $m < 8$ as long as q is not torsion. It is likely that finiteness is not necessary; for instance, the argument of [27] can very likely be modified to deal with the case that $S_{\rho;q;C}$ has components of the form $\tilde{A}_n$. (For small n, the question for affine root systems can be translated via our derived equivalence to one for finite root systems but on a commutative surface.) If we denote the nonisomorphic constituents of $\Phi(\mathcal{O}_{e_m}(-1))$ by $S_0,\ldots,S_n$, with $S_0 \cong \mathcal{O}_e(d)$, then we can work entirely inside the (Serre) subcategory of extensions of the S_i. The algebra $\operatorname{Ext}^*(\oplus_i S_i)$ depends only on the intersection numbers $c_1(S_i)\cdots c_1(S_j)$, and can be recognized as the Koszul dual of the "partial preprojective algebra" Λ [17] of the corresponding quiver. It thus seems very likely that the category generated by the S_i is isomorphic to the category of finite modules over Λ, and thus depends only on the combinatorics of the situation.

Remark. If we denote the pure braid group of the root system $\mathcal{S}_{\rho;q;C}$ (i.e., the kernel of the map from the braid group to the Weyl group) by $P_{S_{\rho;q;C}}$, then we would expect to have an action of the extension $P_{S_{\rho;q;C}}.W(E_{m+1})$ via derived equivalences, acting on $\operatorname{Pic}(X)$ in the natural way. (This may not be faithful in general, although it follows from [57] that it is faithful whenever every component of $\mathcal{S}_{\rho;q;C}$ is of type A.) There is a mild technicality here, in that reflection functors do not act in quite the right way on K_0. Luckily, if the corresponding -2-sheaf is $\mathcal{O}_\alpha(d_\alpha - 1)$, then we need simply twist by $d_\alpha\alpha$ before or after reflecting to fix everything. The case $\alpha = e_{m-1} - e_m$ is particularly interesting. If $d = d_{e_{m-1}-e_m} < 0$, then we find

$$\begin{aligned} R_{e_{m-1}-e_m}(\alpha_m^*\mathcal{O}_{e_{m-1}}(-1))(de_{m-1} - de_m) &\cong \mathcal{O}_{e_m}(-1) \\ R_{e_{m-1}-e_m}(\mathcal{O}_{e_m}(-1))(de_{m-1} - de_m) &\cong \alpha_m^!\mathcal{O}_{e_{m-1}}(-1), \end{aligned} \tag{12.53}$$

while every other object in the exceptional collection is taken to the corresponding object. We may thus think of the combined operation as being a mutation of the exceptional collection, with something similar applying to the inverse operation when $d \geq 0$, or to the other types of simple reflections. This observation will be key to understanding birational geometry when C degenerates, see [44].

Corollary 12.21. *Suppose $m \neq 8$ and the hypotheses of Theorem 12.20 hold for $X_{\rho;q;C}$. Then any surface $X_{\rho';q';C'}$ derived equivalent to $X_{\rho;q;C}$ is Morita equivalent to $X_{\rho;q;C}$.*

Proof. Indeed, reflection functors do not change the surface, and the other generators are abelian equivalences. □

Remark. For $m < 8$, one can omit the hypothesis that q is non-torsion here, as it is straightforward to show in that case that the known derived equivalences surject onto the subgroup of $\operatorname{Aut}(K_0)$ that preserves $\chi(_,_)$, and thus account for all possible destination surfaces. Similarly, for $m = 8$, it suffices to also throw in generators of $\mathrm{SL}_2(\mathbb{Z})$. For $m > 8$, however, the stabilizer of the pairing is almost certainly larger than $\operatorname{Pic}(X) \rtimes W(E_{m+1})$.

We are also interested in abelian equivalences. Here, things are much better behaved.

Proposition 12.22. *Suppose* $\Phi : \operatorname{coh} X_{\rho;q;C} \cong \operatorname{coh} X_{\rho';q';C'}$ *is an abelian equivalence. Then there is an automorphism* $\phi : C \cong C'$ *taking* q *to* q' *and* ρ *to an element of the same* $\operatorname{Pic}(X) \rtimes W(E_{m+1})$ *orbit as* ρ'*. Moreover, if* $\phi = \mathrm{id}$*, then the equivalence is (up to natural isomorphism) the composition of a twist with an admissible element* $w \in W(E_{m+1})$*.*

Proof. Here Φ must induce an equivalence $\operatorname{coh} C \cong \operatorname{coh} C'$, which must therefore preserve the class of a point. It follows that the equivalence preserves rank, and thus Φ in particular takes -2-sheaves to -2-sheaves. But then we can argue as in Theorem 12.20; since Φ preserves the property of having no maps to or from -2-sheaves, we can skip all steps involving reflection functors, so never leave the realm of abelian equivalences. □

Corollary 12.23. *The automorphism group of* $\operatorname{coh} X_{\rho;q;C}$ *can be identified with the subgroup of* $\operatorname{Aut}(C) \times \operatorname{Pic}(X) \rtimes W(E_{m+1})$ *that preserves the parameters and the set of effective roots. The induced action on* $\mathcal{S}pl(X_{\rho;q;C})$ *preserves the Poisson structure up to a root of unity which is trivial on the intersection of* $\operatorname{Aut}(\operatorname{coh} X_{\rho;q;C})$ *with* $\operatorname{Pic}^0(C) \times \operatorname{Pic}(X) \rtimes W(E_{m+1})$*.*

Proof. For the statement about the Poisson structure, recall that to determine the Poisson structure, we needed to specify a nonzero holomorphic differential on C. The group $\operatorname{Pic}(X) \rtimes W(E_{m+1})$ acts trivially on ω_C, as does $\operatorname{Pic}^0(C) \subset \operatorname{Aut}(C)$. □

Remark. For *derived* automorphisms, we instead (subject to our usual finiteness assumption) obtain a group of the form

$$P_{S_{\rho;q;C}}.G \tag{12.54}$$

where G is the subgroup of $\operatorname{Aut}(C).\operatorname{Pic}(X) \rtimes W(E_{m+1})$ that preserves ρ and q. (The only thing to be shown is that any element of $W(E_{m+1})$ has a unique representation as the product of an element of $W(S_{\rho;q;C})$ and an admissible reflection.) There is, of course, the possibility that the pure braid group does not act faithfully.

Note that if q is preserved by a non-translation automorphism of C, then either q is 2-torsion or q is 3-torsion and $j(C) = 0$, and thus this does not arise in truly noncommutative cases. It is, however, useful to include *contravariant* equivalences. Since ad inverts q, we can combine it with any hyperelliptic involution on C to obtain a contravariant equivalence between surfaces with the same value of q, and this will again be a Poisson map.

We thus obtain a Poisson action of the appropriate subgroup of $(\operatorname{Pic}^0(C) \rtimes \mathbb{Z}/2\mathbb{Z}) \times (\operatorname{Pic}(X) \rtimes W(E_{m+1})$ (preserving ρ and the space of effective roots) on the moduli space of simple sheaves. Since the subspace fixed by a Poisson automorphism is again Poisson, this gives a way to construct even more Poisson moduli spaces.

Of particular interest are cases in which the new moduli spaces are themselves examples of the original construction. As an example, consider a surface with parameters

$$\eta, y_0, y_1, \sqrt{q}\eta/y_1, y_2, \sqrt{q}y_2, y_3, \sqrt{q}y_3, \dots \tag{12.55}$$

Translating by $\sqrt{q}$ turns this into

$$q\eta, \sqrt{q}y_0, \sqrt{q}y_1, q\eta/y_1, \sqrt{q}y_2, qy_2, \sqrt{q}y_3, q/y_3, \dots \tag{12.56}$$

Reflecting in the commuting reflections $f - e_1 - e_2$, $e_3 - e_4$, $e_5 - e_6$, etc., all of which are generically effective, then gives

$$q\eta, y_0, y_1, \sqrt{q}\eta/y_1, qy_2, \sqrt{q}y_2, qy_3, \sqrt{q}y_3, \dots \tag{12.57}$$

which becomes the original surface when we twist by $s+f-e_3-e_5-\cdots$ Consider a 1-dimensional sheaf which is isomorphic to its image under this isomorphism. Using the prescription of [47] for how twisting and reflections act on difference equations (as justified below), we find that the equation can be factored in the form

$$v(qz) = C(\sqrt{q}z)C(z)v(z) \tag{12.58}$$

and thus we can express any solution in terms of solutions of the $\sqrt{q}$-difference equations $v(\sqrt{q}z) = \pm C(z)v(z)$. (Note here that C can have additional singularities at the ramification points of the relevant involutions.)

The elliptic Riemann-Hilbert correspondence considered below suggests considering the same configuration with $\sqrt{q}$ replaced by $\sqrt{p}$; this is similarly related to a moduli space of $\sqrt{p}$-elliptic difference equations. There are presumably also analogues corresponding to the map from $q^{1/r}$-difference equations to q-difference equations or from $p^{1/r}$-elliptic difference equations to p-elliptic difference equations.

Another case meriting further investigation is the composition

$$X_{\rho;q;C} \to X_{\rho;1/q;C} \to X_{\iota_\eta\circ\rho;q;C} \to X_{\rho;q;C} \tag{12.59}$$

where ι_η is the hyperelliptic involution corresponding to η and the last isomorphism comes from the central element of D_m (assuming m even). The difference equations preserved by this contravariant equivalence would have the form $v(qz) = A(z)v(z)$ with $A(z)CA(z)^t \propto C$ for some constant matrix C, so that in particular we would have components of the moduli spaces in which A is an orthogonal or symplectic similitude. It follows from the above considerations that the corresponding moduli spaces are Poisson, but of course one would like to know their dimensions (and especially if there are rigid or 2-dimensional cases). More generally, we could hope to consider moduli spaces of equations in which A takes values in some fixed reductive group.

13 Sheaves as difference equations

13.1 Spaces of meromorphic solutions

As the reader may recall, our original motivation for studying these noncommutative surfaces was to understand moduli spaces of (symmetric) elliptic difference equations and maps between them. Though we used an association between sheaves on X_m and difference equations to motivate our constructions, there remain several loose ends. In particular, there are a number of analytic questions to be addressed regarding this association. As a result, we restrict our attention in this section to the case $C \cong \mathbb{C}^*/\langle p\rangle$.

Since our interpretation of sheaves as difference equations relies on the functor $\mathrm{Hom}(_, \mathcal{M}er)$, we first need to understand this functor better. In particular, it is natural to ask whether this functor is exact, i.e., whether the sheaf $\mathcal{M}er$ is injective. We first ask a somewhat simpler question: If $\phi : \mathcal{O}_X(D_1) \to \mathcal{O}_X(D_2)$ is a nonzero morphism of line bundles, is $\mathrm{Hom}(\phi, \mathcal{M}er)$ surjective? Since this acts as a (symmetric elliptic) difference operator on the space of symmetric meromorphic functions in the form $\mathrm{Hom}(\mathcal{O}_X(D_2), \mathcal{M}er)$, we see that we must understand the solution spaces of inhomogeneous symmetric difference equations.

If we do not impose any symmetry condition on the operator, the homogeneous case reduces easily to results of [42].

Proposition 13.1. *Let $0 < |q| < 1$ and let $\mathcal{D}$ be a q-difference operator $\sum_{m_1\le k\le m_2} c_k(z)T^k$ such that each $c_k(z)$ is meromorphic on $\mathbb{C}^*$ and $c_{m_1}(z), c_{m_2}(z) \neq 0$. Then $\ker(\mathcal{D})$ is a vector space of dimension $m_2 - m_1$ over the space of q-elliptic functions.*

Proof. Since $\ker(\mathcal{D}) = \ker(T^{-m_1}\mathcal{D})$, we may as well assume that $m_1 = 0$, $m_2 = m$. We can then rewrite the equation $\mathcal{D}f = 0$ as a system of first-order equations in variables $v_k(z) = T^{k-1}f(z)$:

$$Tv_k(z) = v_{k+1}(z), \quad 1 \leq k < m$$
$$Tv_m(z) = \sum_{1 \leq k < m} \frac{-c_{k-1}(z)}{c_m(z)} v_k(z)$$

Since the corresponding shift matrix is invertible, we may apply [42, Thm. 3] to conclude that there is an m-dimensional space of solutions over the field of T-invariant functions (i.e., the space of q-elliptic functions). □

Corollary 13.2. *Let $0 < |q| < 1$ and let $\mathcal{D} = \sum_{m_1 \leq k \leq m_2} c_k(z)T^k$ be a meromorphic q-difference operator such that $c_{m_1+m_2-k}(z) = c_k(\eta/q^{m_1+m_2}z)$ and $c_{m_1}(z) \neq 0$. Then the space of functions f such that $\mathcal{D}f = 0$ and $f(\eta/z) = f(z)$ is a vector space of dimension $m_2 - m_1$ over the field of η-symmetric q-elliptic functions.*

Proof. Again, we may reduce to the case $m_1 = 0$, $m_2 = m$, and have already observed that $\ker(\mathcal{D})$ is an m-dimensional vector space over the field of q-elliptic functions. Since the field of q-elliptic functions is Galois over the field of η-symmetric q-elliptic functions, and the Galois group preserves $\ker(\mathcal{D})$ (as follows from the symmetry conditions), the claim follows from Galois theory. □

Remark. In principle, this should also follow from [42, Thm. 3], in the case that the group acting is the infinite dihedral group rather than the infinite cyclic group; the translation from straight-line equation to cocycle is somewhat unpleasant, however.

To understand inhomogeneous operators, it remains only to prove existence of solutions. For unconstrained operators, this again follows immediately from [42].

Lemma 13.3. *Let $\mathcal{D}$ be a nonzero meromorphic q-difference operator. Then for any meromorphic function g on $\mathbb{C}^*$, there is a meromorphic function f on $\mathbb{C}^*$ such that $\mathcal{D}f = g$.*

Proof. Reduce to a matrix equation as above and apply [42, Thm. 4]. □

Corollary 13.4. *Let $\mathcal{D} = \sum_{m_1 \leq k \leq m_2} c_k(z)T^k$ be a nonzero meromorphic q-difference operator such that $c_{m_1+m_2-k}(z) = c_k(\eta/q^{m_1+m_2}z)$. Then for any meromorphic function g on $\mathbb{C}^*$ such that $g(\eta/q^{m_1+m_2}z) = g(z)$, there is a meromorphic function f on $\mathbb{C}^*$ such that $\mathcal{D}f = g$ and $f(\eta/z) = f(z)$.*

Proof. Apply the lemma to find a function $h(z)$ such that $\mathcal{D}h = g$, and observe that $\hat{h} := h(\eta/z)$ also satisfies $\mathcal{D}\hat{h} = g$, so that $f := (h + \hat{h})/2$ is the desired solution. □

In terms of our surfaces, this has the following rephrasing as the desired special case of injectivity of $\mathcal{M}er$.

Corollary 13.5. *Let D_1, D_2 be divisor classes and $\phi : \mathcal{O}_X(D_1) \to \mathcal{O}_X(D_2)$ a nonzero morphism. Then* $\mathrm{Hom}(\phi, \mathcal{M}er)$ *is surjective.*

We next deal with blowups.

Lemma 13.6. *Let $\mathcal{M}er_m$ denote the sheaf $\mathcal{M}er$ on X_m. Then $\mathcal{M}er_m \cong \alpha_m^* \mathcal{M}er_{m-1}$ for $m > 0$.*

Proof. We first note that $\alpha_{m*}\mathcal{M}er_m \cong \mathcal{M}er_{m-1}$ (by inspection since $\mathcal{M}er_m$ is saturated and its restriction to $\alpha_m^* \operatorname{Pic}(X_{m-1})$ is precisely $\mathcal{M}er_{m-1}$). Moreover, since the generators of $[e_m]$ and $[f - e_m]$ are invertible as difference operators, we find that the natural map

$$\alpha_m^*\alpha_{m*}\mathcal{M}er_m \to \mathcal{M}er_m \tag{13.1}$$

is surjective; indeed, the image is equal to $\mathcal{M}er_m$ as a module, not just as a sheaf. Moreover, invertibility of $[e_m]$ also implies that

$$\operatorname{Hom}(\mathcal{O}_{e_m}(-1), \mathcal{M}er_m) = 0, \tag{13.2}$$

and thus the natural map

$$\mathcal{M}er_m \to \alpha_m^!\alpha_{m*}\mathcal{M}er_m \cong \alpha_m^!\mathcal{M}er_{m-1} \tag{13.3}$$

is injective. Since $T_{\mathcal{M}er}$ is an isomorphism, we may conjugate by θ to obtain an injective map $\mathcal{M}er_m \to \alpha_m^*\mathcal{M}er_{m-1}$ with direct image the identification $\alpha_{m*}\mathcal{M}er_m \cong \mathcal{M}er_{m-1}$. Composing with the surjection $\alpha_m^*\mathcal{M}er_{m-1} \to \mathcal{M}er_m$ gives a map $\mathcal{M}er_m \to \mathcal{M}er_m$ agreeing with the identity on $\alpha_m^* \operatorname{Pic}(X_{m-1})$, and thus (since elements of those degrees generate the module) agreeing with the identity everywhere. It follows that $\mathcal{M}er_m$ is a direct summand of $\alpha_m^*\mathcal{M}er_{m-1}$. The other direct summand is necessarily a power of $\mathcal{O}_{e_m}(-1)$, and since there are no maps $\alpha_m^*\mathcal{M}er_{m-1} \to \mathcal{O}_{e_m}(-1)$, we find that the map is an isomorphism as required. □

Lemma 13.7. *If $T_M = 0$, then $R\operatorname{Hom}(M, \mathcal{M}er) = 0$.*

Proof. This follows immediately from the fact that $T_{\mathcal{M}er}$ is an isomorphism. □

In particular, this implies that $\mathcal{M}er_{m-1}$ is α_m^*-acyclic, as it has no 0-dimensional subsheaf.

Proposition 13.8. *For any divisor class D, $\operatorname{Ext}^p(\mathcal{O}_X(D), \mathcal{M}er) = 0$ for $p > 0$.*

Proof. Since $\mathcal{M}er$ is translation-invariant, it suffices to consider the case $D = 0$. For $m > 0$, we find

$$\begin{aligned} R\operatorname{Hom}(\mathcal{O}_{X_m}, \mathcal{M}er_{X_m}) &\cong R\operatorname{Hom}(\mathcal{O}_{X_m}, L\alpha_m^!\mathcal{M}er_{X_{m-1}}) \\ &\cong R\operatorname{Hom}(R\alpha_{m*}\mathcal{O}_{X_m}, \mathcal{M}er_{X_{m-1}}) \\ &\cong R\operatorname{Hom}(\mathcal{O}_{X_{m-1}}, \mathcal{M}er_{X_{m-1}}), \end{aligned} \tag{13.4}$$

and thus we reduce by induction to the case $m = 0$. Now, as above, we find that if (ρ_*, ρ^*) are the natural functors between $\operatorname{qcoh}(X_0)$ and $\operatorname{qcoh}(\mathbb{P}^1)$, then

$$\rho^*\rho_*\mathcal{M}er \to \mathcal{M}er \tag{13.5}$$

is surjective (though only as a map of sheaves), since the submodule of $\mathcal{M}er$ generated by elements of degree in $\mathbb{Z}f$ contains all elements of degree D with $D \cdot f \geq 0$. It follows that $\mathcal{M}er$ is acyclic for ρ_*, and thus

$$R\operatorname{Hom}(\mathcal{O}_X, \mathcal{M}er) \cong R\operatorname{Hom}(\rho^*\mathcal{O}_{\mathbb{P}^1}, \mathcal{M}er) \cong R\operatorname{Hom}(\mathcal{O}_{\mathbb{P}^1}, R\rho_*\mathcal{M}er) \tag{13.6}$$

Since $R\rho_*\mathcal{M}er$ is the sheaf on $\mathbb{P}^1$ induced by a vector space over $\mathbb{C}(\mathbb{P}^1)$, it is injective, and thus the claim follows. □

Call a coherent sheaf "vertical" if its Chern class is orthogonal to f.

Lemma 13.9. *Suppose M is a vertical sheaf. Then $R\operatorname{Hom}(M, \mathcal{M}er) = 0$.*

Proof. Since the claim holds for 0-dimensional sheaves, we may reduce to considering irreducible 1-dimensional sheaves. Since $R^1\alpha_{m*}M$ is 0-dimensional, we may similarly reduce to the case $m = 0$, so that $c_1(M) \propto f$. Then Lemma 11.40 implies $c_1(M) = f$. Twisting by a suitable multiple of s then lets us ensure $\chi(M) = 1$, and by the proof of Lemma 11.40, this implies $M \cong \mathcal{O}_X/\mathcal{O}_X(-f)$. Any nonzero map $\mathcal{O}_X(-f) \to \mathcal{O}_X$ becomes an isomorphism under $\operatorname{Hom}(_, \mathcal{M}er)$, and thus the quotient satisfies $R\operatorname{Hom}(M, \mathcal{M}er) = 0$ as required. □

Lemma 13.10. *If M is a quotient of $\mathcal{O}_X(D)$ for some D, then $\operatorname{Ext}^p(M, \mathcal{M}er) = 0$ for $p > 0$.*

Proof. If $M = \mathcal{O}_X(D)$, then the result is already known, and thus we may assume $\operatorname{rank}(M) = 0$. We also note that if $M = \operatorname{coker}\phi$ with $\phi : \mathcal{O}_X(D_1) \to \mathcal{O}_X(D_2)$, then the claim follows from surjectivity of $\operatorname{Hom}(\phi, \mathcal{M}er)$ together with acyclicity of $\mathcal{O}_X(D)$ for $\operatorname{Hom}(_, \mathcal{M}er)$.

More generally, we may use the translation invariance of $\mathcal{M}er$ to arrange that $D = c_1(M)$. Now, consider the short exact sequence

$$0 \to N \to \mathcal{O}_X(c_1(M)) \to M \to 0. \tag{13.7}$$

The sheaf N has numerical invariants $(1, 0, \chi(N))$ where

$$\chi(N) = \frac{c_1(M) \cdot (c_1(M) + C_m)}{2} + 1 - \chi(M). \tag{13.8}$$

In other words, N corresponds to a point in $\operatorname{Hilb}^n(X)$, where $n = 1 - \chi(N)$. By the proof of Lemma 11.43, there is an injective morphism $\mathcal{O}_X(-nf) \to N$. Consider the quotient short exact sequence

$$0 \to N/\mathcal{O}_X(-nf) \to \mathcal{O}_X(c_1(M))/\mathcal{O}_X(-nf) \to M \to 0. \tag{13.9}$$

Since $c_1(N/\mathcal{O}_X(-nf)) = nf$, the first term is annihilated by $R\operatorname{Hom}(_, \mathcal{M}er)$, and thus

$$\operatorname{Ext}^p(M, \mathcal{M}er) \cong \operatorname{Ext}^p(\mathcal{O}_X(c_1(M))/\mathcal{O}_X(-nf), \mathcal{M}er) = 0 \tag{13.10}$$

for $p > 0$. □

Remark. The key idea above was the association of a straight-line equation to a quotient of a line bundle in such a way as to not change the space of solutions. If $\dim \operatorname{Ext}^1(\mathcal{O}_X(-nf), N) = d > 0$, then the association in the proof is not unique, but this can be fixed. Indeed, we find that $\dim \operatorname{Hom}(\mathcal{O}_X(-(n-d)f), N) = 1$ and any map $\mathcal{O}_X(-nf) \to N$ factors through the unique (mod scalars) map $\mathcal{O}_X(-(n-d)f) \to N$. We thus find that the straight-line equation arising in the proof is unique once we remove common factors of the coefficients.

Corollary 13.11. *Let M be a pure 1-dimensional coherent sheaf, and let $\phi : \mathcal{O}_X \to M$ be a morphism. Then for any morphism $\psi : M \to \mathcal{M}er$ the symmetric meromorphic function corresponding to $\psi \circ \phi$ satisfies a symmetric difference equation of order $\leq c_1(M) \cdot f$.*

Given the above, it is straightforward to obtain the following.

Proposition 13.12. *The sheaf $\mathcal{M}er$ is injective.*

Proof. Let M be any coherent sheaf, and let $\phi : \mathcal{O}_X(D) \to M$ be any nonzero morphism. Then $\operatorname{im}\phi$ is a quotient of a line bundle, so that $\operatorname{Ext}^p(\operatorname{im}\phi, \mathcal{M}er) = 0$ for $p > 0$, while $\operatorname{Ext}^p(\operatorname{coker}\phi, \mathcal{M}er) = 0$ by induction on the rank and first Chern class. This implies that $\operatorname{Ext}^p(M, \mathcal{M}er) = 0$ for $p > 0$ as required. Since $\operatorname{Hom}(_, \mathcal{M}er)$ respects short exact sequences of coherent sheaves, it is exact on $\operatorname{qcoh} X_m$ as required. □

We can moreover use similar ideas to understand the space of solutions more precisely. Here there is a technical issue to be considered: although $\mathrm{Hom}(\mathcal{O}_X(-D), \mathcal{M}er)$ is naturally a vector space over the field of $q^{D\cdot f+1}\eta$-symmetric q-elliptic functions, there is no such structure for $\mathrm{Hom}(M, \mathcal{M}er)$ in general. However, if we fix a surjection $\mathcal{O}_X(-D)^n \to \mathcal{M}er$, then this allows us to view $\mathrm{Hom}(M, \mathcal{M}er)$ as a subspace $\mathrm{Hom}(\mathcal{O}_X(-D), \mathcal{M}er)^n$, which *is* a vector space over the field of symmetric elliptic functions.

Theorem 13.13. *Let M be a 1-dimensional sheaf transverse to the anticanonical curve. Then for all sufficiently ample D, the natural surjection*

$$\mathcal{O}_X(-D) \otimes_k \mathrm{Hom}(\mathcal{O}_X(-D), M), \mathcal{O}_X(-D) \to M \tag{13.11}$$

identifies $\mathrm{Hom}(M, \mathcal{M}er)$ *with a vector space of dimension $c_1(M) \cdot f$ over the field of $q^{D\cdot f+1}\eta$-symmetric q-elliptic functions.*

Proof. This property is preserved under extensions, so we may reduce to the case that M is irreducible. Then without loss of generality, M is generated by a global section $\mathcal{O}_X \to M$. Fix such a global section, with kernel I. Then for sufficiently ample D, both I and M are $\mathcal{O}_X(-D)$-globally generated and acyclic, so in particular there is a surjection

$$\mathrm{Hom}(\mathcal{O}_X(-D), \mathcal{O}_X) \to \mathrm{Hom}(\mathcal{O}_X(-D), M). \tag{13.12}$$

In particular, it follows that given any morphism $\phi : M \to \mathcal{M}er$, we can compute the composition $\phi_D : \mathcal{O}_X(-D) \otimes_k \mathrm{Hom}(\mathcal{O}_X(-D), M) \to \mathcal{M}er$ from the composition $g : \mathcal{O}_X \to \mathcal{M}er$ as the images of suitable difference operators. Conversely, given ϕ_D, we can determine the image of g under *all* operators of degree $-D$, which by saturation of $\mathcal{M}er$ allows us to determine g. We thus reduce to showing that the image of $\Gamma(M)$ in $\Gamma(\mathcal{M}er)$ is a space of the requisite dimension.

Twisting by $c_1(M)$, we reduce as above to a short exact sequence

$$0 \to N \to \mathcal{O}_X(c_1(M)) \to M \to 0, \tag{13.13}$$

and may further choose an injection $\mathcal{O}_X(-nf) \to N$. Then up to a 0-dimensional subsheaf, $N/\mathcal{O}_X(-nf)$ is transverse to C, and thus so is $\mathcal{O}_X(c_1(M))/\mathcal{O}_X(-nf)$. It follows that the difference operator corresponding to the map $\mathcal{O}_X(-nf) \to \mathcal{O}_X(c_1(M))$ has nonzero constant term, and thus that the corresponding solution space has dimension equal to the degree of the operator, namely

$$(c_1(M) + nf) \cdot f = c_1(M) \cdot f. \tag{13.14}$$

The claim then follows via the identification

$$\mathrm{Hom}(M, \mathcal{M}er) \cong \mathrm{Hom}(\mathcal{O}_X(c_1(M))/\mathcal{O}_X(-nf), \mathcal{M}er). \tag{13.15}$$

□

Remark. Note that if $\mathrm{rank}(M) > 0$, then we have an injective map $\mathcal{O}_X(D) \to M$ (there is a surjection from a sum of line bundles to M, and at least one of those bundles must have image of positive rank), so that $\mathrm{Hom}(M, \mathcal{M}er)$ has the infinite-dimensional space $\mathrm{Hom}(\mathcal{O}_X(D), \mathcal{M}er)$ as a quotient. Similarly, if M is 1-dimensional but not transverse to C, then we may replace it with the maximal subsheaf transverse to C without changing the corresponding space of solutions.

Though the connection to straight-line forms was convenient for understanding inhomogeneous equations, we should also make the connection to equations in matrix form precise. The construction is essentially that of [47]. Clearly, we need only consider the case $m = 0$.

Proposition 13.14. *Let M be a sheaf on X_0 such that*

$$R^1\rho_* M = R^1\rho_* M(-s) = 0, \tag{13.16}$$

and such that $\rho_* M$ $\rho_* M(-s)$ are vector bundles on $\mathbb{P}^1$. Then there is a short exact sequence

$$0 \to \rho^* V(-s) \to \rho^* W \to M \to 0 \tag{13.17}$$

with vector bundles V, W.

Proof. Since $D^b \operatorname{coh}(X_0)$ is generated by $\mathcal{O}_X$, $\mathcal{O}_X(-f)$, $\mathcal{O}_X(-s-f)$, $\mathcal{O}_X(-s-2f)$, there is a complex representing M in which the terms are direct sums of these line bundles. In particular, we find that there is a distinguished triangle

$$\rho^* V(-s) \to \rho^* W \to M \to \tag{13.18}$$

with $V, W \in D^b \operatorname{coh}(\mathbb{P}^1)$; more precisely, $\rho^* W$ is the complex in which we omit the terms $\mathcal{O}_X(-s-f)$ and $\mathcal{O}_X(-s-2f)$, and similarly for $\rho^* V$. Applying $R\rho_*$ and $R\rho_*(_(-s))$ gives the desired result. □

Remark. Note that (again following [47]) we have isomorphisms

$$\begin{aligned}
\operatorname{Hom}(R^1\rho_* M, \mathcal{O}_x) &\cong \operatorname{Hom}(M, \theta\rho^* \mathcal{O}_x) \\
\operatorname{Hom}(R^1\rho_* M(-s), \mathcal{O}_x) &\cong \operatorname{Hom}(M, \theta\rho^* \mathcal{O}_x(s)) \\
\operatorname{Hom}(\mathcal{O}_x, \rho_* M) &\cong \operatorname{Hom}(\rho^* \mathcal{O}_x, M) \\
\operatorname{Hom}(\mathcal{O}_x, \rho_* M(-s)) &\cong \operatorname{Hom}(\rho^* \mathcal{O}_x(s), M).
\end{aligned}$$

In particular, $R^1\rho_* M$ and $R^1\rho_* M(-s)$ vanish unless M has a map to some sheaf of the form $\mathcal{O}_f(-1)$, while $\rho_* M$ and $\rho_* M(-s)$ are vector bundles unless M has a map from a sheaf of the form $\mathcal{O}_f$. In particular if M is pure 1-dimensional, then when the hypotheses are violated, we can replace M by a sheaf with smaller $c_1(M) \cdot (s+f)$ and with the same space of solutions.

Given such a presentation, we can identify $\operatorname{Hom}(M, \mathcal{M}er)$ with a subspace of $\operatorname{Hom}(\rho^* W, \mathcal{M}er)$, which in turn is equal to a space of vectors (of length $\operatorname{rank}(W)$) of $q\eta$-symmetric meromorphic functions. (Here we use the fact that $\operatorname{Hom}(\mathcal{O}_X(ef), \mathcal{M}er)$ is independent of e; when we consider "holomorphic" solutions below, this will no longer be the case.)

If M is pure 1-dimensional, then Chern class considerations tell us that $\operatorname{rank}(V) = \operatorname{rank}(W)$, and we thus find that $\operatorname{Hom}(M, \mathcal{M}er)$ can be identified with the space of solutions of

$$\begin{aligned}
B(\eta/z)^t w(\eta/z) &= B(z)^t w(z) \\
w(q\eta/z) &= w(z)
\end{aligned}$$

where B corresponds to the map $\rho^* V(-s) \to \rho^* W$. More precisely, if we write

$$\begin{aligned}
V &\cong \bigoplus_{1 \le i \le n} \mathcal{O}_{\mathbb{P}^1}(d_i), \\
W &\cong \bigoplus_{1 \le i \le n} \mathcal{O}_{\mathbb{P}^1}(e_i),
\end{aligned}$$

then $\hat{B} \in \operatorname{Hom}(\rho^* V(-s), \rho^* W)$ corresponds to a matrix of difference operators, such that

$$\hat{B}_{ij} \in \mathcal{S}'(-e_i f, s - d_j f). \tag{13.19}$$

Taking the analytic form of the category, we see that the corresponding operator has the form $\frac{z}{\theta_p(z^2/\eta)}(B(z) - B(\eta/z)T)$, where B is a holomorphic matrix such that

$$B_{ij}(pz) = (-qx_0/z)(\eta/pz^2)^{-d_j}(q\eta/pz^2)^{e_i}B_{ij}(z) \tag{13.20}$$

and thus the solution w must satisfy

$$\sum_i B_{ij}(z)w_i(z) - B_{ij}(\eta/z)w_i(qz) = 0. \tag{13.21}$$

Since $w_i(q\eta/z) = w_i(z)$ and thus $w_i(qz) = w_i(\eta/z)$, this is easy to put in the above form. Note that B will be invertible precisely when T_M is injective.

It also follows from the above proof that the association of V and W to M is functorial. As a result, if $\phi : M \to N$ is a morphism between two sheaves that correspond in this way to matrix difference equations, then there are corresponding morphisms between vector bundles on $\mathbb{P}^1$, and thus there are matrices C, D with $C(\eta/z) = C(z)$, $D(q\eta/z) = D(z)$ such that

$$B_N(z)C(z) = D(z)B_M(z). \tag{13.22}$$

We find that C and D are invertible iff the kernel and cokernel of ϕ have Chern class a multiple of f (or, on some blowup, orthogonal to f). Indeed, consider the commutative diagram of short exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \rho^*V_M(-s) & \xrightarrow{B_M} & \rho^*W_M & \longrightarrow & M & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \rho^*C(-s)} & & \downarrow{\scriptstyle \rho^*D} & & \downarrow{\scriptstyle \phi} & & \\ 0 & \longrightarrow & \rho^*V_N(-s) & \xrightarrow{B_N} & \rho^*W_N & \longrightarrow & N & \longrightarrow & 0. \end{array} \tag{13.23}$$

Now, C and D are invertible iff they are injective with 0-dimensional cokernel. Thus if C and D are invertible, then $\ker\phi$ and $\operatorname{coker}\phi$ are the kernel and cokernel of a map between vertical sheaves. In the other direction, $\rho_*\ker\phi$ would be a 0-dimensional sheaf contained in ρ_*M, and thus $\rho_*\ker\phi = 0$. But this implies that $\rho_*M \subset \rho_*N$, so that D and thus C are injective. Moreover, the cokernel of D has a filtration with subquotients $R^1\rho_*\ker\phi$ and $\rho_*\operatorname{coker}\phi$, both of which are 0-dimensional, so that $\operatorname{coker}\rho^*D$ is vertical, making D invertible. Since we have a four-term exact sequence in which all but one term is vertical, the remaining term $\operatorname{coker}\rho^*C$ must also be vertical, so that C is also invertible.

Note that the corresponding identification between the solution spaces is also an immediate consequence of the fact that $c_1(N/M) \cdot f = 0$, so that $R\operatorname{Hom}(N/M, \mathcal{M}er) = 0$.

Of particular interest is the case that M and N are pseudo-twists (i.e., that N is the direct image of a twist of the minimal lift of M by some e_i, or vice versa). In that case, one finds that the relation between the two equations is precisely as predicted in [47] (apart from the effects of the scalar gauge transformation if we work with the elliptic form of the category), with precisely the same argument as in the commutative case. (The same applies to the other atomic operations considered there.)

To deal with more general gauge/isomonodromy relations, we define an equivalence relation $\sim_h$ on pure 1-dimensional sheaves on X_0 or X_0' to be the weakest equivalence relation such that $M \sim_h N$ if there is a morphism $\phi : M \to N$ with vertical kernel and cokernel.

Proposition 13.15. *Suppose M and N are sheaves with no vertical subsheaf and no pure vertical quotient sheaf. If $M \sim_h N$, then M and N are comparable in the sense of Proposition 11.22.*

Proof. For any coherent sheaf M, let GM denote the quotient of M by its maximal vertical subsheaf. This operation is functorial, and $M \sim_h GM$ via the natural morphism $M \to GM$. Moreover, given a morphism $\phi : M \to N$ with vertical kernel and cokernel, the snake lemma applied to

$$\begin{array}{ccccccccc} 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & GM & \longrightarrow & 0 \\ & & \downarrow & & \phi\downarrow & & G\phi\downarrow & & \\ 0 & \longrightarrow & N' & \longrightarrow & N & \longrightarrow & GN & \longrightarrow & 0 \end{array} \tag{13.24}$$

gives an exact sequence in which $\ker G\phi$ is sandwiched between two vertical sheaves, so is vertical. Since GM has no map from a vertical sheaf, $G\phi$ is injective. Applying $R^1 \operatorname{ad} GR^1 \operatorname{ad}$ (i.e., taking the largest horizontal subsheaf with *pure* vertical kernel) gives a morphism which is still injective but now has 0-dimensional cokernel.

In other words, if we take any sequence of pure 1-dimensional sheaves in which consecutive sheaves are related by such morphisms in either direction, then applying $R^1 \operatorname{ad} GR^1 \operatorname{ad} G$ to the sequence gives a sequence of *comparable* sheaves. Since the original sheaves M, N have $R^1 \operatorname{ad} GR^1 \operatorname{ad} GM \cong M$ and $R^1 \operatorname{ad} GR^1 \operatorname{ad} GN \cong N$, the claim follows. □

If M is a sheaf corresponding to a difference equation, then we can use the operations of [47] to produce a sheaf with no maps to or from vertical sheaves; indeed, if M has a map to or from some sheaf of the form $\mathcal{O}_f(d)$, then performing suitable pseudo-twists in one of the points of intersection turns the sheaf into a subsheaf $\mathcal{O}_f(-1)$ or a quotient $\mathcal{O}_f$, both of which are automatically removed in the translation between sheaves and equations. It follows that if $M \sim_h N$, then we can apply those operations to obtain new sheaves with no vertical sub- or quotient sheaves, and thus apply the Proposition to conclude that the new sheaves are comparable. But then we may apply Proposition 11.22 to conclude that the resulting sheaves are themselves related by a sequence of pseudo-twists. In other words, if the equations corresponding to M and N are related by a gauge/isomonodromy equivalence, then we can make the two equations agree via some sequence of these canonical operations.

13.2 Sheaves of holomorphic solutions

In [50], the term "isomonodromy" was justified via a notion of "weak monodromy": given a difference equation

$$v(qz) = A(z)v(z), \tag{13.25}$$

with $A(pz) = A(z)$, any symmetric fundamental matrix $M(z)$ necessarily [19] satisfies an equation of the form

$$M(pz) = M(z)\hat{A}(z) \tag{13.26}$$

where $\hat{A}(z)$ is q-elliptic. Of course, $M(z)$ is not unique, and thus $\hat{A}(z)$ is only determined up to a suitable equivalence relation (corresponding to taking the quotient of the category of coherent sheaves on X_0' by the subcategory of 0-dimensional sheaves). This suggests that there should be a relation between the corresponding noncommutative surfaces.

In [30], an analogous construction was given in the case of a discrete connection on a vector bundle on $\mathbb{C}^*/\langle p \rangle$ (satisfying certain genericity conditions that unfortunately exclude pullbacks from $\mathbb{P}^1$). In contrast to the trivial bundle case, Krichever's construction gives a unique choice of $\hat{A}(z)$. A careful consideration of that construction reveals the source of this rigidification: one is looking for a fundamental solution with significant constraints on the singularities of the solutions. The argument of [42] for the existence of meromorphic solutions reduces to the study of a particular

sheaf of solutions, and it turns out that we can adapt Praagman's construction to our purposes. In particular, we will find that the symmetric solutions satisfying suitable constraints on their singularities can be naturally identified with the space of global sections of a certain vector bundle on the quotient of $\mathbb{C}^*/\langle q\rangle$ by the involution, and we can use this identification to define a suitable analogue of a fundamental matrix.

As in the case of meromorphic solutions, it will be convenient to understand this sheaf via descent, and thus first consider the case of a non-symmetric q-difference equation. Consider an equation

$$v(qz) = A(z)v(z) \tag{13.27}$$

where $A(z) \in \mathrm{GL}_n(\mathrm{Mer}(\mathbb{C}^*))$. To define a sheaf, we must make some choices about where the solution can have singularities; with this in mind, we choose a (setwise!) section $\sigma : \mathbb{C}^*/\langle q\rangle \to \mathbb{C}^*$ of the natural quotient map $\pi_q : \mathbb{C}^* \to \mathbb{C}^*/\langle q\rangle$.

Definition 13.1. The *sheaf of σ-holomorphic solutions* of $v(qz) = A(z)v(z)$, denoted $\mathcal{S}_{A,\sigma}$, is the analytic sheaf on $\mathbb{C}^*/\langle q\rangle$ defined by taking $\Gamma(U;\mathcal{S}_{A,\sigma})$ to be the space of vectors $v(z) \in \mathrm{Mer}(\pi_q^{-1}U)^n$ which satisfy $v(qz) = A(z)v(z)$ and are holomorphic on the image of σ.

Although in principle σ could be arbitrarily badly behaved, the corresponding sheaf is relatively insensitive to σ. Define a partial ordering $\prec_A$ on sections by $\sigma \prec_A \sigma'$ if at each point $x \in \mathbb{C}^*/\langle q\rangle$, one of the following occurs:

- If $\sigma'(x) = q^k\sigma(x)$ for $k \geq 0$, then $A(q^{k-1}z)$,...,$A(z)$ are holomorphic at $\sigma(x)$.
- If $\sigma'(x) = q^{-k}\sigma(x)$ for $k \geq 0$, then $A(q^{-1}z)^{-1}$,...,$A(q^{-k}z)^{-1}$ are holomorphic at $\sigma(x)$.

Lemma 13.16. *If $\sigma \prec_A \sigma'$, then $\mathcal{S}_{A,\sigma} \subset \mathcal{S}_{A,\sigma'}$.*

Proof. We need to show $\Gamma(U;\mathcal{S}_{A,\sigma}) \subset \Gamma(U;\mathcal{S}_{A,\sigma'})$, or equivalently that every local solution which is holomorphic at σ is also holomorphic at σ'. If $\sigma'(x) = q^k\sigma(x)$ for $k \geq 0$, then $v(z)$ is holomorphic at $\sigma'(x)$ iff

$$v(q^k z) = A(q^{k-1}z)\cdots A(z)v(z) \tag{13.28}$$

is holomorphic at $\sigma(x)$, which follows from the holomorphy of $A(z)$,...,$A(q^{k-1}z)$ there. Similarly, if $\sigma'(x) = q^{-k}\sigma(x)$ for $k \geq 0$, then $v(x)$ is holomorphic at $\sigma'(x)$ iff

$$v(q^{-k}z) = A(q^{-k}z)^{-1}A(q^{1-k}z)^{-1}\cdots A(z/q)^{-1}v(z) \tag{13.29}$$

is holomorphic at $\sigma(x)$. □

Remark. The Lemma clearly would continue to hold if we had merely insisted that $A(q^{k-1}z)\cdots A(z)$ or $A(q^{-k}z)^{-1}\cdots A(z/q)^{-1}$ (as appropriate) were holomorphic in defining the partial ordering. Most of the time, this does not change the partial ordering, and when it does the change tends to be subtle, as it is related to the presence of apparent singularities.

We then define an equivalence relation $\sigma \sim_A \sigma'$ iff both $\sigma \prec_A \sigma'$ and $\sigma' \prec_A \sigma$; it follows immediately from the Lemma that equivalent sections give rise to equal sheaves of solutions. Note that if $A(z)$ is holomorphic on $\pi_q^{-1}(x)$ then we may replace $\sigma(x)$ by any other point in its orbit without changing the equivalence class of σ. It follows that the equivalence class of σ is determined by its value on the images of the singular points of the equation. Moreover, if $\sigma(x)$ is not a singularity of the equation, then replacing $\sigma(x)$ by $q\sigma(x)$ gives a $\sim_A$-equivalent section.

This gives rise to the following alternate description of the data. For each singularity x of the equation, call it "σ-small" if $\log_q(x/\sigma(\pi_q(x))) \geq 0$ and "σ-large" if $\log_q(x/\sigma(\pi_q(x))) < 0$.

Lemma 13.17. *For any q-difference equation $v(qz) = A(z)v(z)$ with $|q| < 1$ with singular points Σ_A, the $\sim_A$-equivalence classes of sections of π_q are in natural bijection with the decompositions*

$$\Sigma_A = \Sigma_A^0 \cup \Sigma_A^\infty \tag{13.30}$$

satisfying $q^{\mathbb{N}}\Sigma_A^0 \cap \Sigma_A^\infty = \emptyset$.

Proof. Given any section σ, we may take Σ_A^0 to be the set of σ-small elements and Σ_A^∞ to be the set of σ-large elements, and it follows from the Lemma that two sections are $\sim_A$-equivalent iff they induce the same decomposition. Conversely, given such a decomposition, we may define a section as follows:

(a) If $\pi_q^{-1}(x) \cap \Sigma_A^0 \neq \emptyset$, let $\sigma(x)$ be the largest element of the intersection.

(b) If $\pi_q^{-1}(x) \cap \Sigma_A^0 = \emptyset$ but $\pi_q^{-1}(x) \cap \Sigma_A^\infty \neq \emptyset$, let $\sigma(x)$ be q^{-1} times the smallest element of the intersection.

(c) If $\pi_q^{-1}(x) \cap \Sigma = 0$, then let $\sigma(x)$ be the unique representative in the annulus $|q| < z \leq 1$.

It is easy to see in this case that Σ_A^0 is precisely the set of σ-small singularities. □

For $r \in \mathbb{R}^*$, let σ_r be the section with image $|q|r < |z| \leq r$. Note that a singularity x is σ_r-small iff $|x| \leq r$ and σ_r-large iff $|x| > r$.

Lemma 13.18. *For any section σ which is bounded away from 0 and ∞, there is a section $\sigma'_r \sim_A \sigma$ which agrees with σ_r on the complement of a finite set.*

Proof. Since σ is bounded away from 0 and ∞, it follows that the set of σ-small elements is bounded away from ∞ and the set of σ-large elements is bounded away from 0. In particular, it follows that there are only finitely many σ-small singularities which are σ_r-large, and only finitely many σ-large singularities which are σ_r-small. Thus if we change σ_r to agree with σ on those orbits, we obtain $\sigma'_r \sim_A \sigma$ as required. □

Proposition 13.19. *Suppose σ is a section of π_q which is bounded away from 0 and ∞. Then $\mathcal{S}_{A,\sigma}$ is a vector bundle of rank n on $\mathbb{C}^*/\langle q \rangle$.*

Proof. Choose a small annulus $r_1 < |z| < r_2$ on which both $A(z)$ and $A(z)^{-1}$ are holomorphic. Fix $r \in (r_1, r_2)$, and let σ'_r be a section as given by the Lemma. Moreover, for any other $r' \in (r_1, r_2)$, there is a section $\sigma'_{r'}$ agreeing with σ'_r on the subannulus $qr_2 \leq |z| \leq r_1$, agreeing with $\sigma_{r'}$ if $r_1 < |z| \leq r'$ or $qr' \leq |z| < qr_2$, and again determining the same sheaf.

Now, define a collection of open subsets of $\mathbb{C}^*$ as follows. Let $x_1, \ldots, x_l$ be the points where $\sigma'_r \neq \sigma_r$ (and thus $\sigma'_{r'} \neq \sigma_{r'}$). Then for each i there is a neighborhood U_i of x_i and a section σ_i agreeing with σ outside U_i such that $\sigma_i(U_i)$ is the component of $\pi^{-q}(U_i)$ containing $\sigma'_r(x_i)$. (Indeed, we simply need to choose U_i small enough to satisfy the covering space property and to not meet the orbit of any other singularity in the relevant bounding annulus.)

Define two more open subsets U_{l+1} and U_{l+2} of $\mathbb{C}^*/\langle q \rangle$ as follows: U_{l+1} is the (image of the) annulus $|q|r < |z| < r$ punctured at $x_1, \ldots, x_l$, and U_{l+2} is the annulus $|q|r' < |z| < r'$ similarly punctured. Also set $\sigma_{l+1} = \sigma'_r$, $\sigma_{l+2} = \sigma'_{r'}$. Clearly the sets $U_1, \ldots, U_{l+2}$ cover $\mathbb{C}^*/\langle q \rangle$.

Each U_i satisfies the covering space property, and thus the space of meromorphic solutions on $\pi_q^{-1}(U_i)$ can be identified with the space of meromorphic vectors on any one of the components. We thus find in particular that

$$\Gamma(U_i; \mathcal{S}_{A,\sigma}) = \Gamma(U_i; \mathcal{S}_{A,\sigma_i}) = \mathbb{A}(\sigma_i(U_i))^n. \tag{13.31}$$

In particular, we find that each $\Gamma(U_i; \mathcal{S}_{A,\sigma_i})$ is a free module of rank n over $\mathbb{A}(\sigma_i(U_i))$, making it a vector bundle as required. □

Corollary 13.20. *Suppose that σ is bounded away from 0 and ∞. Then for any $A \in \mathrm{GL}_n(\mathrm{Mer}(\mathbb{C}^*))$ and $l \in \mathbb{N}$, one has $\mathcal{S}_{\wedge^l A,\sigma} \cong \wedge^l \mathcal{S}_{A,\sigma}$ and $\mathcal{S}_{S^l A,\sigma} \cong S^l \mathcal{S}_{A,\sigma}$. Moreover, for any $B \in \mathrm{GL}_m(\mathrm{Mer}(\mathbb{C}^*))$, one has $\mathcal{S}_{A\otimes B,\sigma} \cong \mathcal{S}_{A,\sigma} \otimes \mathcal{S}_{B,\sigma}$.*

Proof. The first two claims follow immediately from the above construction: the transition matrices between the different trivializations are products of shifts of A, and thus respect *any* endofunctor of the category of vector spaces over $\mathbb{C}$. The remaining claim follows from the fact that $A \otimes B$ is a direct summand of $\wedge^2(A \oplus B)$, since the construction clearly respects direct sums. □

Remark. Two special cases are of particular import: $\wedge^{n-1}$ tells us that the vector bundles $\mathcal{S}_{A^{-t},\sigma}$ and $\mathcal{S}_{A,\sigma}$ are dual, while $\wedge^n$ tells us that $\det \mathcal{S}_{A,\sigma} \cong \mathcal{S}_{\det A,\sigma}$.

Remark. One reason we prefer to use σ rather than the specific decomposition of singularities is that $\det(A)$ and $A \otimes B$ can have fewer singularities than A or $A \oplus B$, which would make the analogous statements more complicated to state.

Corollary 13.21. *Suppose σ, σ' are two sections bounded away from 0 and ∞. Then there is a natural meromorphic isomorphism $\mathcal{S}_{A,\sigma} \cong \mathcal{S}_{A,\sigma'}$, which is biholomorphic away from the orbits containing singularities which are σ-small and σ'-large or σ-large and σ'-small.*

Proof. Indeed, we may assume without loss of generality that σ' agrees with σ away from those finitely many orbits (as the new σ' will be $\sim_A$-equivalent to the original σ'), at which point it follows directly from the construction that the two vector bundles are isomorphic on the relevant open set. □

Corollary 13.22. *Let σ, σ' be sections bounded away from 0 and ∞, with corresponding decompositions $\Sigma_A = \Sigma^0_{A,\sigma} \uplus \Sigma^\infty_{A,\sigma} = \Sigma^0_{A,\sigma'} \uplus \Sigma^\infty_{A,\sigma'}$. Then*

$$\deg(\mathcal{S}_{A,\sigma'}) = \deg(\mathcal{S}_{A,\sigma}) + \sum_{x \in \Sigma^\infty_{A,\sigma'} \cap \Sigma^0_{A,\sigma}} \mathrm{ord}_x(\det A(z)) - \sum_{x \in \Sigma^0_{A,\sigma'} \cap \Sigma^\infty_{A,\sigma}} \mathrm{ord}_x(\det A(z)). \tag{13.32}$$

Proof. We first note that we may compute the degrees of the two vector bundles using their determinants. The resulting first-order equation may have fewer singular points, but since the additional points do not contribute to the sum, this has no effect. We may thus reduce to the case of a first-order equation $v(qz) = a(z)v(z)$. Now, we may write $a(z) = \tilde{a}(z)r(z)$ where $\tilde{a}(z)$ is holomorphic on

$$(\Sigma^0_{A,\sigma'} \cap \Sigma^\infty_{A,\sigma}) \cup (\Sigma^\infty_{A,\sigma'} \cap \Sigma^0_{A,\sigma}), \tag{13.33}$$

and $r(z)$ is a rational function. Since we then find

$$\deg(\mathcal{S}_{a,\sigma'}) - \deg(\mathcal{S}_{a,\sigma}) = \deg(\mathcal{S}_{\tilde{a},\sigma'}) - \deg(\mathcal{S}_{\tilde{a},\sigma}) + \deg(\mathcal{S}_{r,\sigma'}) - \deg(\mathcal{S}_{r,\sigma}), \tag{13.34}$$

and the first two line bundles agree, we thus reduce to the case of a rational function. By factoring the rational function, we may further reduce to the case $v(qz) = (1 - z/u)v(z)$.

If $u \in \Sigma^\infty_{1-z/u,\sigma}$, then the function $(z/u; q)^{-1}$ is holomorphic and nonvanishing at $\sigma(u)$, and thus in that case

$$\mathcal{S}_{1-z/u,\sigma} = (z/u; q)^{-1} \mathcal{S}_{1,\sigma} \cong \mathcal{O}_{\mathbb{C}^*/\langle q \rangle}. \tag{13.35}$$

Similarly, if $u \in \Sigma^0_{1-z/u,\sigma}$, then

$$\mathcal{S}_{1-z/u,\sigma} = (qu/z; q)\mathcal{S}_{-z/u,\sigma}. \tag{13.36}$$

We thus conclude that

$$\deg \mathcal{S}_{1-z/u,\sigma} = \begin{cases} 0 & u \in \Sigma^{\infty}_{A,\sigma} \\ -1 & u \in \Sigma^{0}_{A,\sigma}. \end{cases} \tag{13.37}$$

□

The argument of the Corollary leads to a general approach to computing the degrees of these sheaves. As above, it suffices to consider first-order equations $v(qz) = a(z)v(z)$. At every point in the image of σ, it is easy to write down the germ of a general solution on the corresponding orbit, which in particular determines a minimal order of vanishing (or maximal order of pole) at each point in the orbit. By Weierstrass' theorem, there is a function $f(z)$ with precisely the corresponding divisor. Since both $f(z)$ and $f(z)^{-1}$ are holomorphic on the image of σ, we find

$$\mathcal{S}_{f(z)/f(qz),\sigma} \cong \mathcal{O}_{\mathbb{C}^*/\langle q \rangle}, \tag{13.38}$$

and thus

$$\mathcal{S}_{a(z),\sigma} \cong \mathcal{S}_{a(z)f(z)/f(qz),\sigma}, \tag{13.39}$$

where now $a(z)f(z)/f(qz)$ is holomorphic and nonvanishing on $\mathbb{C}^*$, so the latter line bundle is independent of σ.

It thus remains only to understand the degree for equations without singularities. More generally, suppose only that $a(z)$ is holomorphic and nonvanishing on the circle $|z| = r$. We can then reconstruct the degree of the line bundle $\mathcal{S}_{a,\sigma_r}$ by simply asking how many zeros (with multiplicity, subtracting poles) the typical solution has on the annulus $|q|r \leq |z| \leq r$. This can be computed via contour integral:

$$\begin{aligned} \int_{|z|=r} \frac{v'(z)dz}{2\pi\sqrt{-1}v(z)} - \int_{|z|=|q|r} \frac{v'(z)dz}{2\pi\sqrt{-1}v(z)} &= \int_{|z|=r} \left(\frac{v'(z)}{v(z)} - \frac{qv'(qz)}{v(qz)} \right) \frac{dz}{2\pi\sqrt{-1}} \\ &= -\int_{|z|=r} \frac{a'(z)dz}{2\pi\sqrt{-1}a(z)}. \end{aligned} \tag{13.40}$$

Although this is difficult to compute in general, it does have one particularly nice property: given a holomorphic family of equations $a(z;w)$ such that no fiber has singularities on the circle $|z| = r$, the corresponding integral is continuous in w, and thus constant. More generally, we can replace the circle $|z| = r$ by any Jordan curve C containing 0 and disjoint from qC, with the obvious definition of σ_C.

Note that we can compute $\deg(\mathcal{S}_{a,\sigma})$ for any (bounded) σ by comparison with $\deg(\mathcal{S}_{a,\sigma_r})$; here we note that the correction terms to the contour integral are just residues at the singular points which are on the wrong side of the circle $|z| = r$.

Let us next consider how to descend through an involution $z \mapsto \eta/z$. If we let R_η denote the corresponding operator on functions, we find that

$$R_\eta \mathcal{S}_{A,\sigma} = \mathcal{S}_{A(\eta/qz)^{-1},\eta/\sigma(\eta/z)}. \tag{13.41}$$

We want this to give an equivariant structure on $\mathcal{S}_{A,\sigma}$, and thus need $A(\eta/qz)^{-1} = A(z)$ (the usual cocycle condition on a symmetric equation) and $\sigma \sim_A \eta/\sigma(\eta/z)$. In terms of the large and small singularities of A, this translates to the condition $\Sigma^{\infty}_A = \eta/q\Sigma^0_A$, which is easy to arrange *unless* A is singular at $\pm\sqrt{\eta/q}$. We thus impose as an additional condition that A and its inverse are holomorphic there. We then have an equivariant structure, and may thus consider the sheaf of symmetric sections (a sheaf on the quotient of $\mathbb{C}^*/\langle q \rangle$ by $z \mapsto \eta/z$). If $A(\pm\sqrt{\eta/q}) = 1$, then the equivariant structure is a pullback from $\mathbb{P}^1$, and we obtain the following.

Proposition 13.23. *Under the above conditions, the sheaf of symmetric solutions of $v(qz) = A(z)v(z)$ and holomorphic on the image of σ is a vector bundle on the quotient $\mathbb{P}^1$ of rank n and degree $\frac{1}{2}\deg \mathcal{S}_{A,\sigma}$, and $\mathcal{S}_{A,\sigma}$ is the pullback of this vector bundle to $\mathbb{C}^*/\langle q\rangle$.*

Remark. More generally, if the matrices $A(\pm\sqrt{\eta/q})$ are holomorphic but not necessarily equal to 1, then they are involutions, and we can compute the degree of the sheaf of symmetric solutions as $\frac{1}{2}(\deg \mathcal{S}_{A,\sigma} - r_+ - r_-)$, where r_+ is the multiplicity of -1 as an eigenvalue of $A(\sqrt{\eta/q})$, and similarly for r_-. Note that when A is singular at $\pm\sqrt{\eta/q}$, we could still obtain an equivariant sheaf by taking the sum or intersection of the two sheaves $\mathcal{S}_{A,\sigma}$ and $\mathcal{S}_{A,\eta/\sigma(\eta/z)}$, at the cost of adding an even more complicated local contribution to the degree.

Before we consider the elliptic case, we briefly digress to consider the case $A(z) \in \mathrm{GL}_n(\mathbb{C}(z))$. Since $A(z)$ has only finitely many singularities in $\mathbb{C}^*$, there are two natural choices of section (up to the equivalence relation): σ_0, in which all singularities are σ_0-large, and σ_∞, in which all singularities are σ_∞-small. Then we may associate to the equation $v(qz) = A(z)v(z)$ the two sheaves $\mathcal{S}_{A,\sigma_0}$ and $\mathcal{S}_{A,\sigma_\infty}$, as well as the natural meromorphic map between them (biholomorphic away from the images of the singularities of A). In the special case that A is holomorphic at 0 and ∞ with $A(0) = A(\infty) = 1$, we find that both sheaves are trivial, and the map between sheaves thus corresponds to a matrix of q-elliptic functions, the "Birkhoff monodromy" of the equation [9].

Thus one might expect the general notion of "isomonodromy" deformation to be related to cases in which changing A (but not σ) gives an isomorphic sheaf. If $v(qz) = B(z)v(z)$ is another equation, then $\mathrm{Hom}(\mathcal{S}_{A,\sigma}, \mathcal{S}_{B,\sigma})$ is represented by matrices $C(z)$ satisfying the (gauge) equation

$$C(qz)A(z) = B(z)C(z) \tag{13.42}$$

and holomorphic on the image of σ. Thus the sheaves are isomorphic iff there is such a matrix with $C(z)^{-1}$ also holomorphic on the image of σ, and indeed one then has $\mathcal{S}_{B,\sigma} = C(z)\mathcal{S}_{A,\sigma}$.

Now, in general, any vector bundle on $\mathbb{C}^*/\langle q\rangle$ can be represented in the form $\mathcal{S}_{\mu,\sigma}$ where $\mu(z) \in \mathrm{GL}_n(\mathbb{C}(z) \cap \mathbb{A}(\mathbb{C}^*))$; note that since $\mu(z)$ has no finite singularities, this is independent of σ. (Note that this multiplier μ is not unique!) If we choose such a multiplier for $\mathcal{S}_{A,\sigma}$, then we have $\mathcal{S}_{A,\sigma} \cong \mathcal{S}_{\mu,\sigma}$, and thus there is a matrix M such that both M and M^{-1} are holomorphic on the image of σ and

$$M(qz) = A(z)M(z)\mu(z)^{-1}. \tag{13.43}$$

This matrix is unique up to the choice of μ and up to automorphisms of the vector bundle. For a symmetric equation, we may choose the multiplier μ so that $\mu(\eta/qz)\mu(z) = 1$ (and so that we have the correct behavior at $\pm\sqrt{\eta/q}$) and then insist that M correspond to an isomorphism of *equivariant* sheaves, or equivalently that $M(\eta/z) = M(z)$. Of course, if $\mathcal{S}_{A,\sigma}$ were trivial, this would simply be a fundamental matrix for the equation; in general, this is the most natural analogue of the fundamental matrix when we are trying to control singularities.

Now, suppose we are given an elliptic equation; more precisely, suppose $|p|, |q| < 1$, $\langle p, q\rangle$ has rank 2, and we are given a "p-theta q-difference equation", i.e., an equation of the form $v(qz) = A(z)v(z)$ such that

$$A(pz) = Cz^{-k}\mu_p(qz)^{-t}A(z)\mu_p(z)^t, \tag{13.44}$$

where $\mu_p(z)$ is the multiplier of some vector bundle on $\mathbb{C}^*/\langle p\rangle$ and $C \in \mathbb{C}^*$, $k \in \mathbb{Z}$ are two auxiliary twisting parameters. (Thus A is a meromorphic map $V^* \to T^*V^* \otimes \mathcal{L}$ for some vector bundle V and line bundle L; we have taken the dual of V to make things appropriately contravariant.) Since $\mu_p(z)$, $\mu_p(z)^{-1}$, and Cz^{-k} have no finite singularities, we find that the set of singularities of A

is invariant under the action of $\langle p \rangle$. It is then natural to impose an additional condition on the decomposition $\Sigma_A = \Sigma_A^0 \cup \Sigma_A^\infty$ of the set of singularities, namely that $p\Sigma_A^0 \subset \Sigma_A^0$; after all, px is certainly smaller than x in the naïve sense.

Equivalently, we may observe that the equation has only finitely many singularities in $\mathbb{C}^*/\langle p \rangle$, and thus to decompose the corresponding orbits, it suffices to choose representatives $x_1, \dots, x_N$ in $\mathbb{C}^*$ of the singularities, and let $\sigma_{\vec{x}}$ be any section such that

$$\Sigma^0_{A,\sigma_{\vec{x}}} = \bigcup_{1 \leq i \leq N} p^{1+\mathbb{N}} x_i. \tag{13.45}$$

Of course, such a section does not exist in general, so we impose the condition $x_j \notin p^{\mathbb{N}} q^{\mathbb{N}} p x_i$ for $1 \leq i, j \leq N$, which clearly ensures that the section exists (and is of course unique up to $\sim_A$).

Now, consider the sheaf $\mathcal{S}_{A,\sigma_{\vec{x}}}$. As before, this is a vector bundle on $\mathbb{C}^*/\langle q \rangle$, and we may choose a representation of this vector bundle via the multiplier $\mu_q(z)$, giving a fundamental solution

$$M(qz) = A(z)M(z)\mu_q(z)^{-1} \tag{13.46}$$

(with M, M^{-1} holomorphic on the image of $\sigma_{\vec{x}}$) as above. If we now define

$$\tilde{A}(z) := M(pz)^t \mu_p(z) M(z)^{-t}, \tag{13.47}$$

then

$$M(pz)^t = \tilde{A}(z) M(z)^t \mu_p(z)^{-1}, \tag{13.48}$$

and

$$\tilde{A}(qz) = C z^{-k} \mu_q(pz)^{-t} \tilde{A}(z) \mu_q(z)^t. \tag{13.49}$$

In other words, the equation

$$v(pz) = \tilde{A}(z) v(z) \tag{13.50}$$

is an equation of precisely the form we have been considering, except that p and q have been exchanged.

There is a more canonical construction of $\tilde{A}(z)$ as follows. The operation $\tilde{T}_p : v(z) \mapsto \mu_p(z)^t v(pz)$ takes solutions of $v(qz) = A(z)v(z)$ to solutions of $\tilde{v}(qz) = Cz^{-k}A(z)\tilde{v}(z)$, and thus gives an isomorphism

$$\tilde{T}_p \mathcal{S}_{A,\sigma_{\vec{x}}} = \mathcal{S}_{Cz^{-k}A,\sigma_{\vec{x}/p}} \cong \mathcal{S}_{A,\sigma_{\vec{x}/p}} \otimes \mathcal{L}_{Cz^{-k}} \tag{13.51}$$

Using the natural meromorphic map

$$\mathcal{S}_{A,\sigma_{\vec{x}/p}} \to \mathcal{S}_{A,\sigma_{\vec{x}}} \tag{13.52}$$

gives a meromorphic map

$$\tilde{T}_p \mathcal{S}_{A,\sigma_{\vec{x}}} \to \mathcal{S}_{A,\sigma_{\vec{x}}} \otimes \mathcal{L}_{Cz^{-k}}, \tag{13.53}$$

which is essentially the inverse transpose of $\tilde{A}$. More precisely, if we identify $\mathcal{S}_{A,\sigma_{\vec{x}}}$ with $\mathcal{S}_{\mu_q(z),\sigma_{\vec{x}}}$ using M, then this map is precisely given by $\tilde{A}^{-t}$.

This allows us to quickly identify the singularities of $\tilde{A}$ on $\mathbb{C}^*/\langle q \rangle$: they are precisely the points $x_1, \dots, x_n$ where $\mathcal{S}_{A,\sigma_{\vec{x}}}$ and $\mathcal{S}_{A,\sigma_{\vec{x}/p}}$ fail to be isomorphic. In other words, the above procedure takes a (twisted) p-theta q-difference equation with chosen representatives of the singularities to a (twisted) q-theta p-difference equation with the same singularities.

This *elliptic Riemann-Hilbert correspondence* is in fact functorial, given a suitable notion of morphism between difference equations. If $v(qz) = A(z)v(z)$, $v(qz) = B(z)v(z)$ are two such equations, with the same values of C and k but possibly different $\mu_{p,A}(z)$, $\mu_{p,B}(z)$, a morphism

from the first equation to the second is defined to be a holomorphic linear transformation $\Phi(z)$ satisfying

$$\Phi(qz)^t B(z) = A(z)\Phi(z)^t$$
$$\Phi(pz) = \mu_{p,B}(z)\Phi(z)\mu_{p,A}(z)^{-1}.$$

(I.e., a map between the corresponding vector bundles that takes solutions of the second equation to solutions of the first equation.) Now, if M_A, M_B are fundamental matrices for $v(qz) = B(z)v(z)$ (relative to compatible choices for the representatives of the two equations), then

$$\tilde{\Phi}(z) := M_A(z)^{-1}\Phi(z)^t M_B(z) \tag{13.54}$$

satisfies

$$\tilde{\Phi}(pz)^t \tilde{A}(z) = \tilde{B}(z)\tilde{\Phi}(z)^t$$
$$\tilde{\Phi}(qz) = \mu_{q,A}(z)\tilde{\Phi}(z)\mu_{q,B}(z)^{-1}.$$

Moreover, $\tilde{\Phi}(z)$ is holomorphic on the image of σ, and thus by the second relation is holomorphic everywhere. In other words, $\tilde{\Phi}(z)$ is a morphism from $v(pz) = \tilde{B}(z)v(z)$ to $v(pz) = \tilde{A}(z)v(z)$, giving us a contravariant functor on difference equations.

If we impose the stronger condition $x_i \notin p^{\mathbb{Z}}q^{\mathbb{Z}}x_j$ for $i \neq j$ on our choices of representative, then the chosen representatives are valid for $\tilde{A}$ as well. Moreover, the condition on M is that it is holomorphic on every $p^{\mathbb{Z}}q^{\mathbb{Z}}$-orbit not meeting the singularities of A (or $\tilde{A}$), as well as at every point of the form $p^{1+\mathbb{N}}q^{-\mathbb{N}}x_i$ or $q^{1+\mathbb{N}}p^{-\mathbb{N}}x_i$. These conditions are symmetric between p and q, and thus we find that $M(z)^t$ is a fundamental solution for $\tilde{A}$. In other words, as long as the singularities satisfy the above condition the functor $A \mapsto \tilde{A}$ is an involution!

This has an immediate consequence for isomonodromy transformations. Changing the choice of representatives of the singularities to $p^{m_1}x_1,\ldots,p^{m_N}x_N$ gives us a new equation $\hat{\tilde{A}}$ satisfying

$$C(pz)^t \hat{\tilde{A}}(z) = \tilde{A}(z)C(z)^t, \tag{13.55}$$

where C is the natural meromorphic map between the two bundles of solutions, so is in particular (twisted) elliptic. If we now swap the role of p and q and use the fact that this correspondence is an involution, we find that given any vector $\vec{m}$ of integers, there is a new equation $\hat{A}$ with singularities $q^{m_1}x_1,\ldots,q^{m_N}x_N$, the same "monodromy" equation $\tilde{A}$, and related to A by $\hat{A}(z)C(z) = C(qz)A(z)$ for a suitable twisted elliptic matrix C.

For simple singularities, we can be more precise about this relation. To be precise, suppose x_1 is a simple pole of the equation (simple zeros are analogous; indeed, the elliptic Riemann-Hilbert correspondence respects duality); that is, A has a simple pole at x_1 and $A(z)^{-1}$ is holomorphic at x_1. Then changing $\vec{x}$ to $(x_1/p, x_2, \ldots, x_n)$ has the effect of changing x_1 from a large singularity to a small singularity. We thus find

$$\mathcal{S}_{A,\sigma_{\vec{x}}} \subset \mathcal{S}_{A,\sigma_{x_1/p,x_2,\ldots,x_n}} \tag{13.56}$$

with quotient supported (scheme-theoretically) on x_1. If $C(z)$ is the corresponding map of vector bundles, then the twisted equation satisfies

$$C(pz)^t \hat{\tilde{A}}(z) = \tilde{A}(z)C(z)^t. \tag{13.57}$$

But then $C(z)$ can also be characterized by the requirement that $\tilde{A}(z)C(z)^t$ be holomorphic with holomorphic inverse at x_1. This is precisely the form of the atomic pseudo-twist described in [47]...

We would like to understand the degrees of the corresponding vector bundles. As before, this reduces to the first-order case, and thus to equations $v(qz) = a(z)v(z)$ such that $a(pz) = Cz^{-k}a(z)$. Now, if the chosen representatives of the singularities are $x_1,\dots,x_N$, then we may write

$$a(z) = C'z^{-k'} \prod_{1\le i\le N} \theta_p(z/x_i)^{m_i} \tag{13.58}$$

for suitable parameters C', k', m_i. Since the equation

$$v(qz) = \theta_p(z/x_i)v(z) \tag{13.59}$$

has fundamental solution $\Gamma_{p,q}(z/x_i)$, we conclude that $\mathcal{S}_{a,\sigma_{\vec{x}}}$ has degree k'. Here we note that

$$Cz^{-k} = \frac{a(pz)}{a(z)} = p^{-k'} \prod_{1\le i\le N} (-x_i/z)^{m_i}, \tag{13.60}$$

and thus

$$p^{k'} = C^{-1} \prod_{1\le i\le N} (-x_i)^{m_i}. \tag{13.61}$$

In particular, we can deduce the degree of the bundle of solutions knowing only the multiplier of the equation and the multiplicities of the singularities.

Note also the immediate consequence that if we are given a holomorphic family of p-theta q-difference equations, possibly with varying locations of (representatives of) singularities, varying multiplier matrices, and even varying p and q, then the corresponding family of solution sheaves has constant degree. A key observation is that this implies that the corresponding family of equations $\tilde{A}$ is again a holomorphic family of equations. (More precisely, since it is only determined up to isomorphism, we can only say that it is locally representable by a holomorphic family of equations.) Indeed, if F denotes the total space of the corresponding $\mathbb{C}^*/\langle q\rangle$-bundle, then F is locally projective, and the family of solution sheaves has a natural $\mathcal{O}_F$-module structure. The fibers of this sheaf are locally free and both the rank and the Euler characteristic are constant in the family; it follows as in the proof of [24, Thm. III.9.9] that the sheaf is flat over the base of the family, and therefore locally free on F. The map we used to construct $\tilde{A}$ is then an everywhere-on-the-base defined family of meromorphic maps between such locally free sheaves, and thus gives rise to a holomorphic family of matrices.

Now, suppose we are given a symmetric equation in the form (which we have tweaked slightly from the usual form to make the action of the functor on parameters nicer)

$$B(p\eta/z)^t v(p\eta/z) = B(z)^t v(z)$$
$$v(pq\eta/z) = v(z)$$

where B is a morphism between pullback vector bundles, twisted by a line bundle; that is, B is a holomorphic matrix satisfying

$$B(pz) = Cz^{-k}\mu_{p,W}(z)B(z)\mu_{p,V}(z)^{-1} \tag{13.62}$$

where

$$\mu_{p,V}(z) = \mathrm{diag}(\dots,(\eta/z^2)^{d_i},\dots)$$
$$\mu_{p,W}(z) = \mathrm{diag}(\dots,(q\eta/z^2)^{e_i},\dots).$$

Then $A(z) = B(p\eta/z)^{-t}B(z)^t$ satisfies

$$A(pz) = C^2\eta^{-k}\mu_{p,W}(qz)^{-t}A(z)\mu_{p,W}(z)^t, \tag{13.63}$$

so gives an elliptic equation of the form we have been considering, singular only where $\det B(z)$ or $\det(B(\eta/z))$ vanish.

Let $x_1,\dots,x_N$ be representatives of the distinct points where $\det(B(z))$ vanishes, and let $x_{N+i} = \eta/x_i$. Suppose moreover that no two elements of the set $\{x_1,\dots,x_{2N}\}$ are in the same $p^{\mathbb{Z}}q^{\mathbb{Z}}$-orbit. (Equivalently, in terms of the original roots of $\det(B(z))$, we need $x_i \notin p^{\mathbb{Z}}q^{\mathbb{Z}}x_j$ and that if $\eta/x_i \in p^{\mathbb{Z}}q^{\mathbb{Z}}x_j$, that $x_i = \eta/x_j$.) Then the equivalence class $\sigma_{\vec{x}}$ is well-defined and $\mathcal{S}_{A,\sigma_{\vec{x}}}$ has a natural equivariant structure for the action of $R_{pq\eta/z}$.

The conditions on the singularities ensure that $B(z)$ is invertible at $\pm\sqrt{p\eta}$, so that $A(\pm\sqrt{p\eta}) = 1$ and $\mathcal{S}_{A,\sigma_{\vec{x}}}$ descends to a sheaf on $\mathbb{P}^1$. Writing this sheaf as $\bigoplus_i \mathcal{O}_{\mathbb{P}^1}(f_i)$, we find that we may take

$$\mu_{q,A}(z) = \operatorname{diag}(\dots,(p\eta/z^2)^{f_i},\dots). \tag{13.64}$$

As above, we may now choose an equivariant isomorphism between $\mathcal{S}_{A,\sigma_{\vec{x}}}$ and $\mathcal{S}_{\mu_{q,A},\sigma_{\vec{x}}}$, or equivalently a fundamental solution

$$\begin{aligned} M(qz) &= A(z)M(z)\mu_{q,A}(z)^{-1} \\ M(pq\eta/z) &= M(z). \end{aligned}$$

The corresponding matrix $\tilde{A}(z)$ satisfies (since M is holomorphic and invertible at $\sqrt{q\eta}$)

$$\begin{aligned} \tilde{A}(qz) &= C^2\eta^{-k}\mu_{q,A}(pz)^{-t}\tilde{A}(z)\mu_{q,A}(z)^t \\ \tilde{A}(q\eta/z) &= \tilde{A}(z)^{-1} \\ \tilde{A}(\sqrt{q\eta}) &= 1 \end{aligned}$$

so gives (as expected) a symmetric q-theta p-difference equation. (We can then take a minimal factorization of $\tilde{A}$ as in [47] to recover $\tilde{B}$ if desired, and find that $\tilde{B}$ has at most the singularities of B.) This gives us a correspondence from symmetric p-theta q-difference equations to q-theta p-difference equations, and extends to a contravariant functor as in the nonsymmetric case.

The isomonodromy transformations may be analyzed as in the nonsymmetric case; replacing x_1 by px_1 has the effect of multiplying a simple zero by p and dividing a simple pole by p, and by symmetry the resulting inclusion of solution sheaves is again supported at precisely the corresponding point. We thus find a relation of the form

$$C(pz)^t\hat{\tilde{A}}(z) = \tilde{A}(z)C(z)^t \tag{13.65}$$

where $C(pq\eta/z) = C(z)$ is determined by the requirement that $\tilde{A}(z)C(z)^t$ be holomorphic with holomorphic inverse at η/x_1.

There is another, more subtle operation in the symmetric case: in addition to shifting the singularities by p, we may also shift the symmetry condition. The simplest version of this involves dividing both η and $\vec{x}$ by p. Then we certainly have an inclusion

$$\mathcal{S}_{A;\sigma_{\vec{x}/p,\eta/\vec{x}}} \subset \mathcal{S}_{A;\sigma_{\vec{x},\eta/\vec{x}}} \tag{13.66}$$

with quotient supported at $x_1,\dots,x_N$. This gives rise to a relation

$$\hat{\tilde{A}}(z)C(z)^t = C(pz)^t\tilde{A}(z), \tag{13.67}$$

where $C(z)$ is holomorphic with $\operatorname{ord}_{x_i} \det(C(z)) = \operatorname{ord}_{x_i} \det(A(z)) = \operatorname{ord}_{x_i} \det(B(z))$. Now, this implies that

$$\tilde{A}(z)C(z)^{-t} \tag{13.68}$$

is holomorphic at $x_1,\dots,x_N$, and thus by symmetry that

$$C(q\eta/z)^t \tilde{A}(z)C(z)^{-t} \tag{13.69}$$

is holomorphic on $\mathbb{C}^*/\langle q\rangle$, so constant. Evaluating this constant at $z = \sqrt{q\eta}$ gives

$$\tilde{A}(z) = C(q\eta/z)^t C(z)^{-t}. \tag{13.70}$$

But this is precisely the standard factorization of $\tilde{A}(z)$, and implies

$$\hat{\tilde{A}}(z) = \tilde{B}(pz)^{-t}\tilde{B}(q\eta/z)^t, \tag{13.71}$$

up to automorphism of the corresponding bundle.

13.3 The elliptic Riemann-Hilbert correspondence

Now, let $\mathcal{Q}_!(X) \subset \operatorname{coh} X$ denote the subcategory of sheaves disjoint from the anticanonical sheaf (necessarily pure 1-dimensional sheaves when q is non-torsion). Geometrically, this is a noncommutative analogue of the category of compactly supported sheaves on the quasiprojective surface $X \setminus C$, while analytically it is a category of difference equations with singularities at specified locations.

Now, let $X_{q,p} := X_{\eta,x_0,\dots,x_m;q;p}$ be a noncommutative rational surface, where we suppose that no root of D_m is effective, and for technical reasons that $x_i^2 \notin q^{\mathbb{Z}}$ for $1 \le i \le m$. More precisely, let $\eta,x_0,\dots,x_m,q \in \mathbb{C}^*$ be such that $x_i/x_j \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$ for $1 \le i < j \le m$ and $\eta/x_ix_j \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$ for $1 \le i,j \le m$. (We will similarly replace ρ by a lift $\rho : \operatorname{Pic}(X) \to \mathbb{C}^*$.)

Now, any sheaf in $\mathcal{Q}_!(X_{q,p})$ is the minimal lift of its direct image in X_0, and that direct image cannot have any subquotients with $c_1 \in (1+\mathbb{N})f$, since this would give points of intersection with the anticanonical curve violating our conditions on $X_{q,p}$. We thus find that $\mathcal{Q}_!(X_{q,p})$ can be identified with a subcategory of the category of elliptic q-difference equations of the form (again tweaking things a bit to make the correspondence nicer)

$$\begin{aligned} B(p\eta/z)^t v(p\eta/z) &= B(z)^t v(z) \\ v(pq\eta/z) &= v(z) \end{aligned}$$

where B is a holomorphic matrix such that

$$B_{ij}(pz) = (-pqx_0/z)(q\eta/z^2)^{e_i} B_{ij}(z)(\eta/z^2)^{-d_j}. \tag{13.72}$$

Moreover, the subcategory consists precisely of the equations corresponding to B such that $\det(B)$ vanishes only at $x_1/pq,\dots,x_m/pq \in \mathbb{C}^*/\langle p\rangle$ and such that at each point, $\operatorname{ord}_{x_i/pq} \det(B) = n - \operatorname{rank}(B(x_i/pq))$. (I.e., the "Fuchsian" equations with singularities at $x_1/pq,\dots,x_m/pq$.)

Applying the elliptic Riemann-Hilbert correspondence gives us a new $pq\eta$-symmetric equation, which can again be factored. Taking into account functoriality gives us the following result.

Theorem 13.24. *For any parameters ρ, p, q with $|p|$, $|q| < 1$, $\operatorname{rank}(\langle p,q\rangle) = 2$, there is a contravariant equivalence $\operatorname{Sol}_{\rho;q;p} : \mathcal{Q}_!(X_{\rho;q;p})^{\mathrm{op}} \cong \mathcal{Q}_!(X_{\rho;p;q})$ such that $\operatorname{Sol}_{\rho;p;q}\operatorname{Sol}_{\rho;q;p} \cong \mathrm{id}$, and acting*

on numerical invariants so that $c_1(\mathrm{Sol}_{\rho;q;p}(M)) = c_1(M)$ *and* $p^{\chi(\mathrm{Sol}_{\rho;q;p}(M))}q^{\chi(M)}\rho(c_1(M)) = 1$. *In addition, if* D *is any divisor class, then there is a natural isomorphism*

$$\mathrm{Sol}_{p^{-(D\cdot_)}\rho;q;p}(M) \cong \mathrm{Sol}_{\rho;q;p}(M)(D). \tag{13.73}$$

Moreover, $\mathrm{Sol}_{\rho;q;p}$ *is locally analytic, in the sense that if* M *is a flat holomorphic family of objects, then* $\mathrm{Sol}_{\rho;q;p}M$ *is a locally holomorphic family.*

Proof. First suppose that $X_{\rho;q;p}$ is a surface for which no root of D_m is effective and $\rho(f-2e_i) \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$ for $1 \leq i \leq m$. In that case one may take $\mathrm{Sol}_{\rho;p;q}$ to be the elliptic Riemann-Hilbert correspondence. Note that a priori $\mathrm{Sol}_{\rho;q;p}(M)$ is only defined on an *analytic* vector bundle on $\mathbb{P}^1$, but per GAGA this can be identified with a unique algebraic vector bundle structure. Moreover, if M ranges over a holomorphic family, then we can cover the base by the open subsets on which the solution sheaf is d-regular, allowing us to algebraize $\mathrm{Sol}_{\rho;q;p}(M)$ on such an open subset in a holomorphic way.

It is straightforward to see that $\tilde{A}$ has only simple zeros at the points $x_1/pq,\ldots,x_m/pq$, and thus indeed corresponds to a (unique!) element of $\mathcal{Q}_!(X_{\rho;p,q})$, giving a functor as required. Moreover, the order of the equation is the same, so that $c_1(\mathrm{Sol}(M)) \cdot f = c_1(M) \cdot f$, and the multiplicities of the singularities are the same, so that $c_1(\mathrm{Sol}(M))\cdot e_i = c_1(M)\cdot e_i$. Since $c_1(\mathrm{Sol}(M))\cdot K = c_1(M)\cdot K = 0$, the remaining degree of freedom of the Chern class must agree as well.

The relation between the Euler characteristics then follows from the fact that M and $\mathrm{Sol}(M)$ are disjoint from their respective anticanonical curves. (It could also be recovered from the above general formulas.)

For twisting, we need only verify the claim on a basis of $\mathrm{Pic}(X)$, say $e_1,\ldots,e_m, s-e_1-\cdots-e_m, f$. All but the last of these correspond to the operations already considered above, and per the computations in [47] (and extended to the noncommutative case above) the result indeed agrees with the twist. It thus remains only to show that

$$\mathrm{Sol}_{\eta,x_0/p,x_1,\ldots,x_m;q;p}(M) \cong \mathrm{Sol}_{\eta,x_0,x_1,\ldots,x_m;q;p}(M)(f). \tag{13.74}$$

Since shifting x_0 has the effect of multiplying the twisted matrix $A(z)$ by $p\eta/z^2$, the matrix $\tilde{A}(z)$ is actually unchanged, and simply has a different value of $\mu_q(z)$ (and a corresponding change to C.) In other words, the only effect is to twist the underlying vector bundle by $\mathcal{O}_{\mathbb{P}^1}(1)$, again agreeing with the action of twisting by f.

It remains to prove the result when the above parameter constraints are violated. We observe that for such a functor to exist for $X_{\rho;q;p}$, it suffices for there to be *some* rational pencil R such that (a) for any root r with $R\cdot r = 0$, $\rho(r) \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$, and (b) for any -1-class e with $R\cdot e = 0$, $\rho(R-2e) \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$. Indeed, we can then apply suitable elements of $W(E_{m+1})$ (preserving the set of effective roots) to take R to f, and observe that the resulting surface admits a well-defined Riemann-Hilbert correspondence.

Next, we note that if one blow up an additional point x_{m+1}, then most of the time one has an algebraic isomorphism $\alpha_{m+1}^* : \mathcal{Q}_!(X_m) \cong \mathcal{Q}_!(X_{m+1})$. Indeed, suppose x_{m+1} is chosen to be in the complement of the countably many hypersurfaces $\rho(v) \in p^{\mathbb{Z}}q^{\mathbb{Z}}$ where v ranges over elements of $\Lambda_{E_{m+1}}$ with $v\cdot e_{m+1} \neq 0$. Then any object $M \in \mathcal{Q}_!(X_{m+1})$ has $c_1(M) \in \Lambda_{E_{m+1}}$ and $\rho(c_1(M)) \in p^{\mathbb{Z}}q^{\mathbb{Z}}$, so that $c_1(M)\cdot e_{m+1} = 0$. But this implies $M \cong \alpha_{m+1}^*\alpha_{(m+1)*}M$, giving the desired isomorphism.

Now, starting with our original surface X_m, perform $2m+1$ such blowups, and consider the divisor class

$$R_m := 2^{m-1}s + 2^m f - \sum_{1\leq k\leq m-1} 2^{m-1-k}e_k - \sum_{1\leq k\leq m} 2^{m-k}(e_{m+k}+e_{2m+k}) - e_{3m+1}. \tag{13.75}$$

on X_{3m+1}. If X_m were itself sufficiently general, then R_m would be the class of a rational pencil on X_{3m+1}. Indeed, this is certainly true for $R_0 = s + 2f - e_2 - e_3 - e_4$, while in general if we reflect R_m in $f - e_{m+1} - e_{2m+1}$ then $s - e_1$, the result is in the same S_{3m+1}-orbit as R_{m-1}.

For general X_m, the algorithm for testing whether R_m is a rational pencil involves repeated reflections in roots having negative intersection with R_m, but we easily see that R_m has positive intersection with every root of E_{m+1}. It follows that if r is any root with $r \cdot R_m \neq 0$, then $r \notin \Lambda_{E_{m+1}}$, and thus (by the assumption on our blowups) $\rho(r) \notin p^{\mathbb{Z}} q^{\mathbb{Z}}$. In other words, the procedure never involves reflecting in an effective root, and thus R_m is indeed a pencil on our surface.

It remains to show that R_m satisfies the requirements to induce a Riemann-Hilbert correspondence. That $\rho(r) \notin p^{\mathbb{Z}} q^{\mathbb{Z}}$ for roots orthogonal to R_m follows by the same argument, since again this implies $r \notin \Lambda_{E_m}$. Similarly, if e is a -1-class orthogonal to R_m, then $(R_m - 2e) \cdot e_{3m+1} = 1 + 2(e \cdot e_{3m+1}) \neq 0$, so that $\rho(R_m - 2e)$ is a function of x_{3m+1}. Thus at the cost of excluding a further countable set of values of x_{3m+1}, we may ensure that $\rho(R_m - 2e) \notin p^{\mathbb{Z}} q^{\mathbb{Z}}$ for any such e. It follows that there is indeed a well-defined Riemann-Hilbert correspondence on $\mathcal{Q}_!(X_{3m+1})$ corresponding to R_m. □

Remark. We will see below that the Riemann-Hilbert functor Sol commutes with the actions of simple reflections whenever both surfaces admit Riemann-Hilbert functors. This will then imply that the construction via auxiliary blowups agrees with the straightforward Riemann-Hilbert functor when the latter is defined.

Remark. The parameter constraints actually ensure that any connected holomorphic family M is bounded, and thus we could omit the word "locally" above. However, we will not be using this fact below, and there are variants for which we cannot guarantee global analyticity.

Remark. The elliptic Riemann-Hilbert functor "integrates" twisting by line bundles, in that we can express such a twist as a composition of two elliptic Riemann-Hilbert functors:

$$_(D) \cong \mathrm{Sol}_{q^{-(D\cdot_)}\rho;p;q} \mathrm{Sol}_{\rho;q;p}\,. \tag{13.76}$$

This is, of course, directly analogous to the role of the usual Riemann-Hilbert correspondence in solving the Painlevé VI equation.

It turns out that the condition $\mathrm{rank}(\langle p, q\rangle) = 2$ is not actually necessary for the conclusion to hold. Indeed, when $\langle p, q\rangle$ has rank 1, the quasiprojective surfaces $X_{\rho;q;p}[1/T]$ and $X_{\rho;p;q}[1/T]$ are Azumaya algebras over the same commutative quasiprojective surface (with anticanonical curve $\mathbb{C}^*/\langle p, q\rangle$). Thus the existence of a locally analytic equivalence $\mathcal{Q}_!(X_{\rho;q;p}) \cong \mathcal{Q}_!(X_{\rho;p;q})$ would follow (apart from some mild technical issues) if we could prove that the two Azumaya algebras were Morita equivalent. In fact, we have the following.

Proposition 13.25. *Let X be a commutative rational surface over $\mathbb{C}$ and let C be an anticanonical curve on X. Any holomorphic Azumaya algebra on $X \setminus C$ is trivial.*

Proof. Since the map $\mathrm{Br}(X \setminus C) \to H^2((X \setminus C)^{\mathrm{an}}; \mathcal{O}_X^*)$ is injective, it suffices to show that the latter is trivial. Since $\hat{X} := X \setminus C$ is noncompact, it follows from [56, Prop. 2.1] that $H^2(\hat{X}^{\mathrm{an}}; \mathcal{O}_{\hat{X}}^*)$ is isomorphic to the purely topological group $H^3(\hat{X}; \mathbb{Z})$. We will show that $H^p(\hat{X}; \mathbb{Z}) = 0$ for $p \geq 3$.

Suppose $\pi : X \to Y$ is a birational morphism with exceptional curve e, so that $\pi(C)$ is again anticanonical on Y. Then $\hat{Y}$ is naturally identified with an open subset of $\hat{X}$ with complement $e \setminus (e \cap C) \cong \mathbb{A}^1(\mathbb{C})$. Let U be a tubular neighborhood of the complement. Then U retracts to $\mathbb{A}^1(\mathbb{C})$, so is contractible, while $U \cap \hat{Y} \cong \mathbb{A}^1(\mathbb{C}) \times \mathbb{C}^*$ is homotopy equivalent to S^1. Applying Mayer-Vietoris to $\hat{X} = \hat{Y} \cup U$ thus implies $H^p(\hat{X}; \mathbb{Z}) \cong H^p(\hat{Y}; \mathbb{Z})$ for $p \geq 3$.

We may thus reduce to the case that X is an odd Hirzebruch surface. In that case, we find that if $X \cong F_d$, then $dC - (d-1)F_d$ is an ample divisor with reduced support C. It follows that $\hat{X}$ is affine, and thus a Stein manifold. The claim then follows from [2, Thm. 1] and the universal coefficient theorem. □

Remark. This holds more generally for any Poisson surface which is not symplectic. Again we may reduce to ruled surfaces, and find that $\hat{X}$ is affine, a line bundle over a Riemann surface, or a $\mathbb{C}^*$-bundle over an elliptic curve. In the first two cases, we again find $H^3(\hat{X};\mathbb{Z}) = 0$, while in the third case $H^3(\hat{X};\mathbb{Z}) \cong \mathbb{Z}$. Either way, the torsion subgroup of $H^3(\hat{X};\mathbb{Z})$ is trivial, and the claim follows.

If we view $\mathrm{Sol}(M)$ as the elliptic analogue of monodromy, then we see that we should expect to obtain an isomonodromy transformation from any isomorphism between surfaces $X_{\rho;p,q}$. We have already considered twisting by line bundles, and will consider the action of $W(E_{m+1})$ more carefully below. There are some other isomorphisms, however: it is also true that the surface only depends on the image of p in $\mathbb{C}^*/\langle q\rangle$, and more subtly that the construction of the surface respects isomorphisms between curves and thus respects *modular* transformations.

To understand these latter transformations, replace ρ with a map $\rho : \mathrm{Pic}(X) \to \mathbb{C}$ and p, q by a triple $\omega_1, \omega_2, \omega_3$ of periods. If ω_2/ω_1 is in the upper half-plane, then we may define a surface

$$X_{\rho;\omega_1,\omega_2,\omega_3} := X_{\exp(2\pi\sqrt{-1}\rho/\omega_1);\exp(2\pi\sqrt{-1}\omega_3/\omega_1);\exp(2\pi\sqrt{-1}\omega_2/\omega_1)}. \tag{13.77}$$

To extend this to the case that ω_2/ω_1 is in the lower half-plane, we take the convention that

$$X_{\rho;q;p} := X^{\mathrm{op}}_{1/\rho;1/q;1/p} \tag{13.78}$$

if $|p| > 1$.

Corollary 13.26. *Let $\omega : \mathbb{Z}^3 \to \mathbb{C}$ be a linear map. For any $g \in \mathrm{GL}_3(\mathbb{Z})$ such that both $\langle \omega_1, \omega_2\rangle$ and $\langle (g\omega)_1, (g\omega)_2\rangle$ are lattices, there is a locally analytic equivalence*

$$\mathcal{Q}_!(X_{\rho;\omega}) \cong \mathcal{Q}_!(X_{\det(g)\rho;g\omega}). \tag{13.79}$$

Proof. Note first that $X^{\mathrm{op}}_{\rho;q;p} \cong X_{\rho;1/q;p} \cong X_{1/\rho;q;p}$, from which one concludes that $X_{-\rho;\omega} \cong X^{\mathrm{op}}_{\rho;\omega}$.

If $\langle \omega_1, \omega_2, \omega_3\rangle$ is a lattice, then $X_{\rho;\omega}[1/T]$ is an Azumaya algebra, so is analytically trivial by Proposition 13.25 above, and is thus Morita equivalent to its center Z. In other words, $X_{\rho;\omega}[1/T]$ is the endomorphism ring of an analytic vector bundle V on Z. Now, let Y be any (possibly nonreduced) projective subvariety of $Z \setminus C$. Since Y is projective, $V|_Y$ is isomorphic to a unique *algebraic* vector bundle on Y, and thus induces a Morita equivalence between coherent sheaves on $X_{\rho;\omega}$ supported on Y and coherent sheaves on Z supported on Y. Taking the union over all Y gives the desired equivalence $\mathcal{Q}_!(X_{\rho;\omega}) \cong \mathcal{Q}_!(Z)$. Now, Z is the commutative rational surface with $C = \mathbb{C}/\langle \omega_1, \omega_2, \omega_3\rangle$ and induced parameter map ρ. Since this is GL_3-invariant (and $\mathcal{Q}_!(X_{\rho;1;p})$ is self-dual), the existence of a family of algebraic equivalences follows immediately. That the family is locally analytic follows once we observe that $V|_Y$ is locally algebraizable in flat families.

We now consider the case that $\langle \omega_1, \omega_2, \omega_3\rangle$ has rank 3. Since X is functorial in the curve, the result clearly holds for the subgroup $\mathrm{ASL}_2(\mathbb{Z})$ preserving $\langle \omega_1, \omega_2\rangle$. Moreover, this extends to $\mathrm{AGL}_2(\mathbb{Z})$ since our conventions ensure $X_{\rho;\omega_1,\omega_2,\omega_3} \cong X_{-\rho;-\omega_1,\omega_2,\omega_3}$. Moreover, the elliptic Riemann-Hilbert correspondence gives $\mathcal{Q}_!(X_{\rho;\omega_1,\omega_2,\omega_3}) \cong \mathcal{Q}_!(X_{-\rho;\omega_1,\omega_3,\omega_2})$ whenever both $\langle \omega_1, \omega_2\rangle$ and $\langle \omega_1, \omega_3\rangle$ are lattices. (Note that there are four cases to consider, depending on which half-space contains ω_2/ω_1 and ω_3/ω_1.) The above elements generate $\mathrm{GL}_3(\mathbb{Z})$, so the result follows as long as the image of every rank 2 subgroup is a lattice, or equivalently when ω has dense image.

In the remaining case, the topological closure of the image of ω is either a line (in which case *no* rank 2 subgroup has image a lattice) or a union of cosets of a line. Using the $\mathbb{C}^*$ symmetry, we may assume that the closure consists of those elements of $\mathbb{C}$ with integer imaginary part. In particular, we have a vector $\alpha := \Im(\omega)$, with the same action of $\mathrm{GL}_3(\mathbb{Z})$. The condition that ω_1, ω_2 generate a lattice is then that $(\alpha_1, \alpha_2) \neq 0$. Using the $\mathrm{ASL}_2(\mathbb{Z})$ symmetry, we have a locally analytic equivalence with a surface for which $\alpha_2 = 0$, in which case we may apply the Riemann-Hilbert functor to obtain a surface with $\alpha_3 = 0$. A further application of $\mathrm{ASL}_2(\mathbb{Z})$ then gives a surface with $\alpha_2 = \alpha_3 = 0$, and thus (up to $-1 \in \mathrm{SL}_2(\mathbb{Z})$) $\alpha_1 = -1$.

It thus remains only to show that the claim holds when both surfaces have $\alpha = (-1, 0, 0)$. The stabilizer of this vector is a different copy of $\mathrm{AGL}_2(\mathbb{Z})$, so we need only generate this group. The original $\mathrm{ASL}_2(\mathbb{Z})$ includes operations $\omega_1 \mapsto \omega_1 + \omega_2$ and $\omega_3 \mapsto \omega_3 + \omega_2$, and it is easy to see that these operations together with the Riemann-Hilbert correspondence generate $\mathrm{AGL}_2(\mathbb{Z})$ as required. □

Remark. It is unclear whether the relations of GL_3 (relative to the above implicit presentation) act trivially, or as nontrivial (holomorphic!) autoequivalences of $\mathcal{Q}_!$.

Remark. There is, of course, a nonsymmetric version of this action coming from the nonsymmetric elliptic Riemann-Hilbert correspondence. Now, if A and $\tilde{A}$ are related by the elliptic Riemann-Hilbert correspondence, then the the fundamental matrix exhibits a cohomology between the (nonabelian) 1-cocycles

$$v(qz) = A(z)v(z) \quad v(pz) = \mu_p(z)v(z) \tag{13.80}$$

and

$$v(qz) = \mu_q(z)^{-t}v(z) \quad v(pz) = \tilde{A}(z)^{-t}v(z). \tag{13.81}$$

Replacing z by $\exp(2\pi\sqrt{-1}z/\omega_1)$ gives a pair of cohomologous 1-cocycles for $\langle \omega_1, \omega_2, \omega_3 \rangle$. We thus find more generally that any two equations related by the general action of (the appropriate extension of) $\mathrm{SL}_3(\mathbb{Z})$ correspond to cohomologous 1-cocycles (and similarly in the symmetric case, now on the group $\langle \omega_1, \omega_2, \omega_3 \rangle \rtimes \mathbb{Z}/2\mathbb{Z}$.) Given any relation, the corresponding product of fundamental matrices induces an automorphism of the appropriate equation, and will thus in many cases be forced to be scalar. If this could be made precise, this would give rise to a nonabelian analogue of the results of [21] on the elliptic Gamma function, replacing the latter by the appropriate fundamental matrices.

Remark. Of course, this extends to an action of $\mathrm{Pic}(X)^3 \rtimes \mathrm{GL}_3(\mathbb{Z})$ (modulo autoequivalences): $\mathrm{Pic}(X)^2$ acts by changing ρ to a different representative modulo $\langle \omega_1, \omega_2 \rangle$, while the third $\mathrm{Pic}(X)$ acts by twisting.

It is tempting to combine the above result with our derived equivalences for $m = 8$ to obtain an action of $\mathrm{SL}_4(\mathbb{Z})$. There are some technical issues, however. The first is that in order for $\mathcal{Q}_!(X_{\rho;\omega})$ to be nonempty, there needs to be a nef divisor class with $\rho(D) \in \operatorname{im}\omega$; in the $m = 8$ case, this imposes an integer linear dependence between the four parameters, and thus some points in the $\mathrm{SL}_4(\mathbb{Z})$-orbit do not correspond to surfaces. More seriously, there is a potential issue in showing that $D^b\mathcal{Q}_!(X_{\rho;\omega})$ is preserved by the derived equivalences. The difficulty here is that the natural subcategory of $D^b \operatorname{coh} X$ is the kernel of $|^{\mathbf{L}}_C$, or equivalently the category of bounded complexes with cohomology in $\mathcal{Q}_!$. But not every such complex is represented by a complex in $\mathcal{Q}_!$. Indeed, if we consider a surface X such that the only irreducible object in $\mathcal{Q}_!(X)$ is $\mathcal{O}_r$ for a fixed -2-curve r, then $\mathrm{Ext}^1_X(\mathcal{O}_r, \mathcal{O}_r) = 0$ implies that $D^b\mathcal{Q}_!(X) \cong D^b\mathcal{V}ect$. It follows that $\mathrm{Ext}^2_{\mathcal{Q}_!(X)}(\mathcal{O}_r, \mathcal{O}_r) = 0$, while Poincaré duality tells us that $\mathrm{Ext}^2_X(\mathcal{O}_r, \mathcal{O}_r) \cong \mathbb{C}$.

It turns out, however, that if we impose additional conditions on the surface, we can arrange for the derived equivalence to restrict to an *abelian* equivalence on $\mathcal{Q}_!(X)$, giving the fol-

lowing. For notation, given $\rho : \mathrm{Pic}(X_7) \to \mathbb{C}$ and $\omega : \mathbb{Z}^4 \to \mathbb{C}$, define a surface $Y_{\rho;\omega} := X_{\rho,\rho(2s+3f-e_1-\cdots-e_7)+\omega_4;\omega_1,\omega_2,\omega_3}$.

Theorem 13.27. *Let $\omega : \mathbb{Z}^4 \to \mathbb{C}$ be a linear map with rank 3 image and $\rho : \mathrm{Pic}(X_8) \to \mathbb{C}$ be such that for any root r of E_8, $\rho(r) \notin \operatorname{im}\omega$. Fix $g \in \mathrm{SL}_4(\mathbb{Z})$ and suppose that both $\langle \omega_1, \omega_2\rangle$ and $\langle (g\omega)_1, (g\omega)_2\rangle$ are lattices. Then there is a locally analytic equivalence $\mathcal{Q}_!(Y_{\rho;\omega}) \cong \mathcal{Q}_!(Y_{\rho;g\omega})$.*

Proof. Let β be a primitive element of $\ker\omega$. Since $\langle \omega_1, \omega_2\rangle$ is a lattice, it is in particular rank 2, implying $(\beta_3, \beta_4) \neq 0$. We can then apply the $\mathrm{GL}_3(\mathbb{Z})$ action to ensure $\beta_1 = \beta_2 = 0$ and thus $\gcd(\beta_3, \beta_4) = 1$.

Now, let $g \in \mathrm{SL}_4$ be an element of the SL_2 acting on the last two coordinates such that $(g^t\beta)_4 = 0$, $(g^t\beta)_3 = 1$. The proof of Theorem 11.62 then gives a derived equivalence $D^b \operatorname{coh} Y_{\rho;g\omega} \to D^b \operatorname{coh} Y_{\rho;\omega}$. Now, every object in $\mathcal{Q}_!(Y_{\rho;g\omega})$ is an extension of structure sheaves of points, and thus this derived equivalence takes $\mathcal{Q}_!(Y_{\rho;g\omega}) \to \mathcal{Q}_!(Y_{\rho;\omega})$. Since every irreducible object in $Y_{\rho;\omega}$ is in the image of this functor, it follows that the inverse derived equivalence also preserves the t-structure, and thus we obtain an equivalence of abelian categories as desired.

We may thus reduce to the case $\beta = (0, 0, 0, 1)$, or by using

$$\mathcal{Q}_!(Y_{\rho;\omega_1,\omega_2,\omega_3,0}) \cong \mathcal{Q}_!(Y_{\rho;\omega_1,\omega_2,0,\omega_3}) \tag{13.82}$$

to the case $\beta = (0, 0, 0, 1)$. The stabilizer of this covector is $\mathrm{AGL}_3(\mathbb{Z})$, but only the $\mathrm{GL}_3(\mathbb{Z})$ quotient acts nontrivially on the parameters, and thus a further application of Corollary 13.26 gives the desired result. □

Remark. Note that essentially the entire content of this Theorem is in the word "analytic". Indeed, without the holomorphic structure, the category $\mathcal{Q}_!(Y_{\rho;\omega_1,\omega_2,0,\omega_3})$ is simply the direct sum of uncountably many copies of the category of sheaves on a surface supported on a single smooth point, and thus any two such categories are *algebraically* equivalent...

Remark. More generally, we have an (almost) action of the group $\mathbb{C}.\Lambda^4_{E_8} \rtimes (W(E_8) \times \mathrm{SL}_4(\mathbb{Z}))$ on these categories (or, more precisely, an extension of this action by a possibly trivial group of holomorphic autoequivalences), where $\mathbb{C}$ acts on ρ by adding a multiple of $(2, 1, 1, 1, 1, 1, 1, 1, 1)$.

One interesting consequence of the elliptic Riemann-Hilbert correspondence is the following.

Corollary 13.28. *There is a biholomorphic map from the moduli space of simple sheaves on $X_{\rho;q;p}$ disjoint from the anticanonical curve to the corresponding moduli space on $X_{\rho;p;q}$, which can be chosen to be respect the symplectic structures up to a scalar multiple.*

Proof. The nontrivial statement is that the map is symplectic (up to a scalar). The difficulty is that the pairing is defined in terms of the Ext^2 groups in $\operatorname{coh} X$, but the functor is defined on $\mathcal{Q}_!$, and we have seen this need not have the same Ext^2 groups. We will thus need to embed our categories in slightly larger categories so that we can represent all of the relevant self-2-extensions inside the category.

The proof of Theorem 13.24 lets us assume without loss of generality that every object $M \in \mathcal{Q}_!(X_{\rho;q;p})$ satisfies $c_1(M) \cdot f > 0$. Now, blow up two more points in such a way that $\rho(f - e_{m+1} - e_{m+2}) = 1$, and let Sol be the functor on $\mathcal{Q}_! X_{m+2}$ constructed in Theorem 13.24. This restricts to an equivalence

$$\alpha^*_{m+2}\alpha^*_{m+1}\mathcal{Q}_!(X_{\rho;q;p}) \cong \alpha^*_{m+2}\alpha^*_{m+1}\mathcal{Q}_!(X_{\rho;p;q}), \tag{13.83}$$

so still induces a biholomorphic map on moduli spaces. But we now have an additional sheaf $R = \mathcal{O}_{f-e_{m+1}-e_{m+2}}(-1)$ with the property that $\mathrm{Ext}^1(\alpha^*_{m+2}\alpha^*_{m+1}M, R) \neq 0$ for all $M \in \mathcal{Q}_!(X_{\rho;q;p})$. For

simple M, let α be an extension of R by $\alpha_{m+2}^*\alpha_{m+1}^*M$, and β an extension of $\alpha_{m+2}^*\alpha_{m+1}^*M$ by R such that the Serre duality pairing between α and β is nonzero. In particular we find that $\mathrm{Ext}^2(R,R) \cong \mathbb{C}$ is generated by the Yoneda product $\alpha\beta$. We can thus represent any nontrivial 2-extension of R by itself as a complex inside $\mathcal{Q}_!(X_{m+2})$, and can then apply Sol to obtain a map $\mathrm{Ext}^2(R,R) \to \mathrm{Ext}^2(\mathrm{Sol}(R),\mathrm{Sol}(R))$, necessarily an isomorphism. Similarly, we can generate $\mathrm{Ext}^2(M,M) \cong \mathbb{C}$ by $\beta\alpha$, obtaining an isomorphism $\mathrm{Ext}^2(M,M) \to \mathrm{Ext}^2(\mathrm{Sol}(M),\mathrm{Sol}(M))$. Moreover, we have

$$\mathrm{tr}(\mathrm{Sol}(\beta\alpha)) = -\mathrm{tr}(\mathrm{Sol}(\alpha\beta)) = -c\,\mathrm{tr}(\alpha\beta) = c\,\mathrm{tr}(\beta\alpha) \tag{13.84}$$

for some nonzero constant c. It follows that the pullback of the symplectic structure on $X_{\rho;p;q}$ is a multiple of the symplectic structure on $X_{\rho;q;p}$ as required. □

Remark. In fact, since we have not specified a differential on $\mathbb{C}^*/\langle p\rangle$, the two symplectic structures are in principle only defined up to a scalar. Of course, there *is* a natural choice of differential, namely $dz/2\pi\sqrt{-1}z$, and presumably the Riemann-Hilbert correspondence is either symplectic or antisymplectic relative to this choice of differential.

Under some additional assumptions, we can arrange to preserve stability. Given an effective divisor such that $D \cdot C_m = 0$, let $\chi(D)$ denote the Euler characteristic of a sheaf disjoint from C with that Chern class; note that this extends to a linear functional on $\rho^{-1}\langle p,q\rangle$.

Proposition 13.29. *Suppose $|p|,|q| < 1$ with $\mathrm{rank}\langle p,q\rangle = 2$, and fix $\rho : \mathrm{Pic}(X_m) \to \mathbb{C}^*/\langle p\rangle$, $D_0 \in \rho^{-1}\langle p,q\rangle$ a nontrivial effective divisor with $D_0 \cdot C_m = 0$. Suppose D_a is an ample divisor such that for any $D \in \rho^{-1}\langle p,q\rangle \cap C_m^\perp$ with D, $D_0 - D$ nontrivial effective, either*

$$\frac{D}{D\cdot D_a} = \frac{D_0}{D_0 \cdot D_a} \quad \text{or} \quad \frac{\chi(D)}{D\cdot D_a} \neq \frac{\chi(D_0)}{D_0\cdot D_a}. \tag{13.85}$$

Then there is a lift of ρ to $\mathrm{Hom}(\mathrm{Pic}(X_m),\mathbb{C}^*)$ such that for any M with $c_1(M) = D_0$, M is D_a-stable iff $\mathrm{Sol}_{\rho;q;p}(M)$ is D_a-stable, and thus $\mathrm{Sol}_{\rho;q;p}$ induces a biholomorphic map between the corresponding stable moduli spaces.

Proof. Define a map $\bar\chi : \rho^{-1}\langle p,q\rangle \cap C_m^\perp \to \mathbb{Q}$ by

$$\bar\chi(D) = \chi(D) - (D\cdot D_a)\frac{\chi(D_0)}{D_0\cdot D_a}. \tag{13.86}$$

This is clearly linear with $\bar\chi(D_0) = 0$, and is thus determined by its restriction to the negative definite lattice

$$\rho^{-1}\langle p,q\rangle \cap \langle D_a, C_m\rangle^\perp. \tag{13.87}$$

It follows that there is an element

$$D' \in \langle D_a, C_m\rangle^\perp \otimes \mathbb{Q} \tag{13.88}$$

such that $\bar\chi(D) = D'\cdot D$.

Now, as M ranges over all sheaves with $c_1(M) = D_0$ and M disjoint from the anticanonical curve, it follows from Lemma 10.21 that there are only finitely many possible Chern classes of proper nontrivial subsheaves of M. Such a sheaf M is stable iff for every proper nontrivial subsheaf N of M, $D'\cdot c_1(N) = \bar\chi(c_1(N)) < 0$.

Now choose any lift ρ to $\mathbb{C}^*$. Since Sol is contravariant and preserves Chern classes, we see that M is stable iff every proper nontrivial *quotient* N of $\mathrm{Sol}(M)$ has $D'\cdot c_1(N) < 0$. But then for $r \gg 0$,

$$\chi(N(-rD')) > \frac{c_1(N)\cdot D_a}{c_1(M)\cdot D_a}\chi(M). \tag{13.89}$$

Since there are only finitely many possible numerical invariants of such N, we may choose r independently of N, and furthermore insist that $rD' \in \mathrm{Pic}(X)$. For such an r, we find that for quotients of $\mathrm{Sol}(M)$, $D' \cdot c_1(N) < 0$ iff $\chi(N(-rD')) > \frac{c_1(N)\cdot D_a}{c_1(M)\cdot D_a}\chi(M)$, and thus conclude that M is stable iff $\mathrm{Sol}(M)(-rD')$ is stable. The claim follows once we absorb the twist by $-rD'$ into the choice of lift of ρ. □

In the $K^2 = 0$ case, this gives the following.

Corollary 13.30. *For arbitrary parameters such that $|p|$, $|q| < 1$ and $\mathrm{rank}\langle p, q\rangle = 2$, the (commutative) quasiprojective surfaces $X_{\eta,x_0,\dots,x_7,(\eta^2 x_0^3/x_1\cdots x_7)q;1;p}[T^{-1}]$ and $X_{\eta,x_0,\dots,x_7,(\eta^2 x_0^3/x_1\cdots x_7)p;1;q}[T^{-1}]$ are biholomorphic.*

Remark. Of course, this, too, extends to an analytic-isomorphism-class-preserving action of $\mathrm{GL}_3(\mathbb{Z})$, or more precisely $\mathbb{C}.\Lambda_{E_8}^3 \rtimes (W(E_8) \times \mathrm{GL}_3(\mathbb{Z}))$. If we remove all of the -2-curves from the surfaces, then we obtain an honest action of this group (extended by analytic automorphisms as usual). Note that the elliptic Painlevé equation appears inside this action as the action of a subgroup $\mathbb{Z} \subset \Lambda_{E_8}$.

This raises the following question: If X, Y are two quasiprojective surfaces obtained by removing the anticanonical curve from a rational surface with $K^2 = 0$, when are they biholomorphic? (And what is the holomorphic automorphism group of X?)

13.4 The Fourier transform

To understand how the action of $W(E_m)$ interacts with the Riemann-Hilbert functor, it will be helpful to give an alternate construction of the functor. In particular, on the (Morita/Mukai) principle that a functor between categories of sheaves should come from a sheaf on the product, we want to look for a sheaf M_{Sol} on $X_{\rho;q;p} \times X_{\rho;p;q}$ such that $\mathrm{Sol}(M) = \mathrm{Hom}(M, M_{\mathrm{Sol}})$ for $M \in \mathcal{Q}_!(X)$. Since $\mathrm{Hom}(\mathcal{O}_{X_{\rho;p;q}}, \mathrm{Sol}(M))$ is the space of suitably holomorphic solutions of the difference equation corresponding to M, this suggests that we try to construct M_{Sol} so that $\mathrm{Hom}(\mathcal{O}_{X_{\rho;p;q}}, M_{\mathrm{Sol}})$ is a subsheaf of $\mathcal{M}er$.

We in particular want to define a space of meromorphic functions that contains every suitably holomorphic solution of our difference equations but has a more analytic definition. Since we are assuming that the equation has only simple singularities, part of the condition is straightforward: at every point, we have a bound on the order of the pole of the solution; since this includes the requirement that the solution be holomorphic where it is required to be holomorphic, this will certainly contain the functions we require. Unfortunately, the resulting space (even with the assumption of symmetry) is too big; the problem is that our operators have a natural tendency to introduce additional singularities at points with $z^2 \in p^{\mathbb{Z}}q^{\mathbb{Z}}\eta$, and thus tend to end up in a larger space. We thus need to impose additional conditions at those points.

A similar issue arises with the Fourier transformation. Recall that we had to define the generalized Fourier transformation as a purely formal transformation, despite the fact that we can associate to it a perfectly well-behaved kernel function. The difficulty there, of course, is that we need to control the singularities of the input to the transformation in order to specify the contour of integration. Again, though, we find that constraining the singularities alone is not enough to give a well-behaved transformation, as the integral itself can introduce additional singularities. However, in that case, we may use the results of [49, §10] to guide us, in particular Lemma 10.4 op. cit. In addition to constraints on singularities, that result has an additional (much stranger) hypothesis, which suggests the following.

Lemma 13.31. *Let $A(z) \in \mathrm{GL}_n(\mathbb{C}^*/\langle p\rangle)$, and suppose that $A(p\eta/z) = A(z)^{-1}$, $A(\pm\sqrt{\eta}) = A(\pm\sqrt{p\eta}) = 1$, and A is holomorphic at every point with $x^2 \in p^{\mathbb{Z}}q^{\mathbb{Z}}\eta$. If v is any $pq\eta$-symmetric solution of $v(qz) = A(z)v(z)$ which is also holomorphic at such points, then*

$$v(p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta}) = v(p^{(1+i)/2}q^{(1-j)/2}\sqrt{\eta}) = v(p^{(1-i)/2}q^{(1-j)/2}\sqrt{\eta}) = v(p^{(1-i)/2}q^{(1+j)/2}\sqrt{\eta}),$$

for all $i, j \in \mathbb{Z}$ and either square root of η.

Proof. The symmetry of v means that it suffices to prove

$$v(p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta}) = v(p^{(1+i)/2}q^{(1-j)/2}\sqrt{\eta}), \tag{13.90}$$

where we may assume $j \geq 0$. We can now write

$$v(p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta}) = A_j v(p^{(1+i)/2}q^{(1-j)/2}\sqrt{\eta}), \tag{13.91}$$

where

$$A_j = A(q^{-1}p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta})A(q^{-2}p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta})\cdots A(q^{-j}p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta}), \tag{13.92}$$

and the hypotheses ensure that the matrices are all defined. We then find

$$\begin{aligned} A_0 &= 1 \\ A_1 &= A(p^{(1+i)/2}\sqrt{\eta}) \\ A_{j+2} &= A(p^{(1+i)/2}q^{(j+1)/2}\sqrt{\eta})A_j A(p^{(1+i)/2}q^{-(j+1)/2}\sqrt{\eta}). \end{aligned}$$

Since

$$\begin{aligned} A(p^{(1+i)/2}\sqrt{\eta}) &= 1, \\ A(p^{(1+i)/2}q^{(j+1)/2}\sqrt{\eta}) &= A(p^{(1+i)/2}q^{-(j+1)/2}\sqrt{\eta})^{-1}, \end{aligned}$$

it follows by induction that $A_j = 1$ for all j, and the result follows. □

This leads us to the following definition.

Definition 13.2. Let $\eta, x_0, \ldots, x_m, p, q$ be parameters in $\mathbb{C}^*$ such that $|p|, |q| < 1$, $\langle p, q\rangle$ has rank 2 and $x_i^2 \notin p^{\mathbb{Z}}q^{\mathbb{Z}}\eta$ for $0 \leq i \leq m$. Then define $\mathcal{M}_{\eta,x_0,\ldots,x_m;p,q}$ to be the space of functions $g(z) \in \mathrm{Mer}(\mathbb{C}^*)$ such that

(a) $g(pq\eta/z) = g(z)$.

(b) $\Gamma_{p,q}(z/pqx_0, \eta/zx_0)\prod_{1\leq i\leq m}(z/x_i, pq\eta/x_i z; p, q)g(z)$ is holomorphic on $\mathbb{C}^*$

(c) For all $i, j \in \mathbb{Z}$ and either square root of η, $g(p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta}) = g(p^{(1+i)/2}q^{(1-j)/2}\sqrt{\eta})$.

Remark. Note that this "strange" condition is indeed *very* strange from an analytic perspective: the points $p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta}$ and $p^{(1+k)/2}q^{(1+l)/2}\sqrt{\eta}$ can be arbitrarily close together (and indeed, the set has an at least 1-dimensional set of limit points) without implying any particular constraint on the distance between $p^{(1+i)/2}q^{(1-j)/2}\sqrt{\eta}$ and $p^{(1+k)/2}q^{(1-l)/2}\sqrt{\eta}$. As a result, it appears likely to be very difficult to answer even very basic questions about the spaces $\mathcal{M}$. Indeed, apart from cases when we can apply Lemma 13.31, it looks hard to even determine if $\mathcal{M}$ is empty!

To deal with the strange condition, the following result will be useful.

Lemma 13.32. *Suppose $u^2 \notin p^{\mathbb{Z}}q^{\mathbb{Z}}\eta$, and let $h(z) := \Gamma_{p,q}(z/pqu, \eta/zu)$. Then $h(pq\eta/z) = h(z)$ and for all $i, j \in \mathbb{Z}$ and either square root of η,*

$$h(p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta}) = (-pqu/\sqrt{\eta})^{ij}h(p^{(1+i)/2}q^{(1-j)/2}\sqrt{\eta}).$$

Proof. We may write

$$\frac{h(p^{(1+i)/2}q^{(1+j)/2}\sqrt{\eta})}{h(p^{(1+i)/2}q^{(1-j)/2}\sqrt{\eta})} = \frac{\Gamma_{p,q}(p^{(-1+i)/2}q^{(-1+j)/2}\sqrt{\eta}/u, p^{(-1-i)/2}q^{(-1-j)/2}\sqrt{\eta}/u)}{\Gamma_{p,q}(p^{(-1+i)/2}q^{(-1-j)/2}\sqrt{\eta}/u, p^{(-1-i)/2}q^{(-1+j)/2}\sqrt{\eta}/u)}. \tag{13.93}$$

Denoting the right-hand side by $F_{ij}(u)$, we find $F_{i0}(u) = F_{0j} = 1$, and can easily simplify

$$F_{i1}(u) = \frac{\theta_p(p^{(-1+i)/2}\sqrt{\eta}/qu)}{\theta_p(p^{(-1-i)/2}\sqrt{\eta}/qu)}, \tag{13.94}$$

which evaluates to $(-pqu/\sqrt{\eta})^i$ as required, using the quasiperiodicity relation for θ_p. We then find

$$\frac{F_{i(j+2)}(u)}{F_{ij}(u)} = F_{i1}(q^{(-1-j)/2}u)F_{i1}(q^{(1+j)/2}u) = (-pqu/\sqrt{\eta})^{2i}, \tag{13.95}$$

from which the claim follows by induction in $|j - 1/2|$. □

Remark. This makes it straightforward to extend the definition of $\mathcal{M}$ to the case $x_0 \in p^{\mathbb{Z}}q^{\mathbb{Z}}$, by imposing the appropriate condition on $\Gamma_{p,q}(z/pqx_0, \eta/zx_0)g(z)$.

Lemma 13.33. *Let $v(qz) = A(z)v(z)$ be the symmetric q-theta equation corresponding to a sheaf in $M \in \mathcal{Q}_!(X_{\eta,x_0,\dots,x_m;q;p})$, where we suppose $x_i^2 \notin p^{\mathbb{Z}}q^{\mathbb{Z}}\eta$ and no root of D_m is effective. Let g be the function arising from a global section of M and a symmetric global section of $\mathcal{S}_{A,\sigma_{\vec{x}}}$. Then*

$$\Gamma_{p,q}(z/pqx_0, \eta/zx_0)^{-1}g(z) \in \mathcal{M}_{\eta,x_0,\dots,x_m;p,q}. \tag{13.96}$$

Proof. Note that the fundamental matrix used to define the elliptic Riemann-Hilbert correspondence induces a pairing between the respective vector bundles, and thus a pairing on their global sections, allowing us to define g. That g is symmetric and has the required singularities is clear. For the strange condition, we first observe that if $x_0^2 \notin p^{\mathbb{Z}}q^{\mathbb{Z}}\eta$, then changing to the algebraic gauge gives an equation to which Lemma 13.31 applies. The function g then satisfies

$$\Gamma_{p,q}(z/pqx_0, \eta/zx_0)^{-1}g(z) = \sum_i h_i(z)v_i(z) \tag{13.97}$$

where $v(z)$ is a solution of the elliptic equation and $h_i(z)$ are symmetric p-elliptic functions with poles controlled by x_0. It then follows easily that $\Gamma_{p,q}(z/pqx_0, \eta/zx_0)^{-1}g$ satisfies the strange condition. □

The following is trivial.

Proposition 13.34. *For $1 \leq i < m$,*

$$\mathcal{M}_{\eta,x_0,\dots,x_i,x_{i+1},\dots,x_m;p,q} = \mathcal{M}_{\eta,x_0,\dots,x_{i+i},x_i,\dots,x_m;p,q} \tag{13.98}$$

Denote this isomorphism by $s_{e_i-e_{i+1}}$. There is also an isomorphism corresponding to reflection in $f - e_1 - e_2$, though this is somewhat more complicated.

Proposition 13.35. *If $g(z) \in \mathcal{M}_{\eta,x_0,\dots,x_m;p,q}$, then*

$$(s_{f-e_1-e_2}g)(z) := \frac{\Gamma_{p,q}(x_1z/\eta, pqx_1/z, z/pqx_0, \eta/zx_0)}{\Gamma_{p,q}(z/x_2, pq\eta/x_2z, zx_1x_2/pqx_0\eta, x_1x_2/zx_0)}g(z)$$
$$\in \mathcal{M}_{\eta,\eta x_0/x_1x_2,\eta/x_2,\eta/x_1,x_3,\dots,x_m;p,q}.$$

Proof. That $s_{f-e_1-e_2}g$ satisfies the symmetric condition is trivial, and the holomorphy condition says that

$$\Gamma_{p,q}(x_1x_2z/pq\eta x_0, x_1x_2/zx_0)(x_1z/\eta, pqx_1/z, x_2z/\eta, pqx_2/z; p, q)$$
$$\times \prod_{3\le i\le m}(z/x_i, pq\eta/x_iz; p, q)(s_{f-e_1-e_2}g)(z) \tag{13.99}$$

is holomorphic, but this is equal to

$$\Gamma_{p,q}(z/pqx_0, \eta/zx_0)\prod_{1\le i\le m}(z/x_i, pq\eta/x_iz; p, q)g(z), \tag{13.100}$$

so that the holomorphy condition is automatic as well. The strange condition is an immediate consequence of Lemma 13.32. □

As the remaining reflection corresponds to the generalized Fourier transformation, we would hope to be able to express it as a contour integral, and indeed we can do this. Define a kernel function (defined as long as $x_0/pqx_1 \notin p^{\mathbb{N}}q^{\mathbb{N}}$)

$$K(z, w; \eta, x_0, x_1; p, q) :=$$
$$\frac{(p;p)(q;q)}{2\Gamma_{p,q}(x_0/x_1)}\frac{\Gamma_{p,q}(z/w, zw/pq\eta, pq\eta x_0/x_1zw, wx_0/zx_1, x_1w/\eta, pqx_1/w, w/pqx_0, \eta/wx_0)}{\Gamma_{p,q}(x_1z/\eta, pqx_0/z, z/pqx_1, \eta x_0/zx_1^2, w^2/pq\eta, pq\eta/w^2)}. \tag{13.101}$$

This comes from the expression for the kernel of the generalized Fourier transformation given above by applying the appropriate gauge transformation and making suitable choices for lifts of the parameters to $\mathbb{C}^*$.

If the parameters satisfy the additional conditions

$$x_1/x_0 \notin p^{\mathbb{N}}q^{\mathbb{N}}$$
$$x_k/pqx_1 \notin p^{\mathbb{N}}q^{\mathbb{N}}, 2 \le k \le m$$
$$x_kx_l/pq\eta \notin p^{\mathbb{N}}q^{\mathbb{N}}, 2 \le k < l \le m,$$

then for generic z, there exists a contour C which contains the (generically distinct!) points $p^{\mathbb{N}}q^{\mathbb{N}}z, p^{\mathbb{N}}q^{\mathbb{N}}pq\eta x_0/x_1z, p^{\mathbb{N}}q^{\mathbb{N}}pqx_1$ and $p^{\mathbb{N}}q^{\mathbb{N}}pq\eta/x_k, 2 \le k \le m$ and excludes their images under $w \mapsto pq\eta/w$. Then for any $g \in \mathcal{M}_{\eta,x_0,x_1,\dots,x_m;p,q}$, we may define

$$(s_{s-e_1}g)(z) = \int_C K(z, w; \eta, x_0, x_1; p, q)g(w)\frac{dw}{2\pi\sqrt{-1}w}. \tag{13.102}$$

Per [49, §10], this extends to a meromorphic function in z.

Proposition 13.36. *If the parameters satisfy the above conditions as well as*

$$\begin{aligned}\eta/x_0x_1 &\notin p^{\mathbb{Z}}q^{\mathbb{Z}}\\ \eta/x_1x_k &\notin p^{\mathbb{Z}}q^{\mathbb{Z}}, 2 \le k \le m\\ x_k/x_l &\notin p^{\mathbb{Z}}q^{\mathbb{Z}}, 2 \le k < l \le m,\end{aligned}$$

then s_{s-e_1} induces a linear transformation

$$\mathcal{M}_{\eta,x_0,\dots,x_m;p,q} \to \mathcal{M}_{\eta x_0/x_1,x_1,x_0,x_2,\dots,x_m;p,q}. \tag{13.103}$$

Proof. We first note (to let us simplify notation) that if we rescale z, w by u, then this has the effect of rescaling η by u^2 and $x_0,\dots,x_m$ by u. It thus suffices to consider the case $\eta = 1/pq$, which we do to make $g \in \mathcal{M}(\eta, x_0, \dots, x_m; p, q)$ satisfy $g(1/z) = g(z)$.

Now, since both the kernel and the constraints on the contour are symmetric under $z \mapsto x_0/x_1z$, it follows immediately that $(s_{s-e_1}g)(x_0/x_1z) = (s_{s-e_1}g)(z)$. Similarly, once we have shown that $s_{s-e_1}g$ is holomorphic at points of the form $p^{i/2}q^{j/2}\sqrt{x_0/x_1}$, the strange condition follows from the fact that

$$K(p^{i/2}q^{j/2}\sqrt{x_0/x_1}, w) = K(p^{i/2}q^{-j/2}\sqrt{x_0/x_1}, w) \tag{13.104}$$

(appropriately modified if $x_0/x_1^3 \in p^{\mathbb{Z}}q^{\mathbb{Z}}$).

It thus remains only to show the holomorphy condition, or equivalently that

$$\begin{aligned}&(pqx_1z, pqx_0/z; p, q) \prod_{2\le i\le m} (z/x_i, x_0/x_1x_iz; p, q)\\ &\quad\times \int_C \frac{\Gamma_{p,q}(zw^{\pm1}, (x_0/x_1z)w^{\pm1}, pqx_1w^{\pm1}, w^{\pm1}/pqx_0)}{\Gamma_{p,q}(w^{\pm2})} g(w)\frac{dw}{2\pi\sqrt{-1}w}\end{aligned} \tag{13.105}$$

analytically continues to the points where the contour is not defined. But this follows immediately from [49, Lem. 10.4]. □

Remark. In particular, the strange condition on g corresponds precisely to the condition that $\Delta(\pm p^{i/2}q^{j/2}) = -\Delta(\pm p^{i/2}q^{-j/2})$, where $\Delta(w) = K(z, w)g(w)$.

Remark. The conditions

$$\begin{aligned}\eta/x_1x_k &\notin p^{\mathbb{Z}}q^{\mathbb{Z}}, 2 \le k \le m\\ x_k/x_l &\notin p^{\mathbb{Z}}q^{\mathbb{Z}}, 2 \le k < l \le m,\end{aligned}$$

are needed to ensure that the integrand has no multiple poles. Presumably the claim still holds without those conditions, as [49, Lem. 10.4] continues to hold in families violating the multiple pole condition, as long as the *generic* integral in the family has only simple poles.

Remark. In the case $m = 8$, a different analytic representation of $W(E_9)$ (acting in the same way on difference equations!) was considered in [53], which in particular constructed solutions by taking eigenfunctions of the operator corresponding to a translation.

We thus see that for every simple reflection in $W(E_{m+1})$, there is a corresponding linear transformation between spaces $\mathcal{M}$, and thus for any word in $W(E_{m+1})$ we obtain such a linear transformation as long as the intermediate spaces are defined.

Theorem 13.37. *Suppose $s_1 \cdots s_k = s'_1 \cdots s'_l$ are two words for the same element of $W(E_{m+1})$. Then the corresponding linear transformations between spaces $\mathcal{M}$ agree whenever both are defined.*

Proof. Clearly, it suffices to prove this for the Coxeter relations. Most of these relations at most multiply the function by a product of elliptic gamma functions, and the corresponding identity reduces to the reflection principle $\Gamma_{p,q}(x, pq/x) = 1$. All but two of the relations involving $s - e_1$ are similarly trivial (as both sides are integral operators with the same kernel), so that we need only prove

$$s_{s-e_1}^2 = 1,$$
$$s_{s-e_1} s_{e_1-e_2} s_{s-e_1} = s_{e_1-e_2} s_{s-e_1} s_{e_1-e_2}.$$

For the braid relation, we need to prove

$$\int_{C_1}\int_{C_2} K(z, y; \eta x_0/x_1, x_1, x_2) K(y, w; \eta, x_0, \dots, x_m; p, q) g(w) \frac{dw}{2\pi\sqrt{-1}w} \frac{dy}{2\pi\sqrt{-1}y}$$
$$= \int_{C_3} K(z, w; \eta, x_0, x_2) g(w) \frac{dw}{2\pi\sqrt{-1}w}, \tag{13.106}$$

as an identity of meromorphic functions in z. For z sufficiently small, we can interchange the integrals on the left, and thus reduce to showing

$$\int_C K(z, y; \eta x_0/x_1, x_1, x_2) K(y, w; \eta, x_0, \dots, x_m; p, q) \frac{dy}{2\pi\sqrt{-1}y} = K(z, w; \eta, x_0, x_2). \tag{13.107}$$

Without loss of generality, we may rescale the parameters so that $\eta = x_1/pqx_0$, at which point the desired identity is simply the elliptic beta integral [60] with parameters

$$(1/x_0, pqx_2, z, x_1/x_2 z, x_0 w/x_1, 1/w).$$

It remains to show that s_{s-e_1} is an involution. We may view $\mathcal{M}_{\eta,x_0,\dots,x_m}$ as a subspace of $\mathcal{M}_{\eta,x_0,x_1,x_0/u,x_2,\dots,x_m}$ for any u, and then apply the above braid relation to conclude that

$$\int\int K(z, y; \eta x_0/x_1, x_1, x_0/u) K(y, w; \eta, x_0, x_1) g(w) = \int K(z, w; \eta, x_0, x_0/u) g(w). \tag{13.108}$$

Since $s_{s-e_1}^2 g$ is the limit of the left-hand side as $u \to 1$, it remains to show that

$$\lim_{u\to 1} \int K(z, w; \eta, x_0, x_0/u) g(w) = g(z). \tag{13.109}$$

The contour condition cannot be satisfied at $u = 1$, and thus we must deform the contour in order to take the limit. The factor $\Gamma_{p,q}(u)^{-1}$ in the kernel vanishes at $u = 1$, and thus ensures that only the residue terms contribute to the limit. In particular, we find (using symmetry)

$$\lim_{u\to 1} \int K(z, w; \eta, x_0, x_0/u) g(w) = 2 \lim_{u\to 1} \operatorname{Res}_{w=z} K(z, w; \eta, x_0, x_0/u) g(w) = g(z) \tag{13.110}$$

as required. □

Remark. Note that the involution property is a variant of [62, Thm. 4.1], with different conditions on the input functions. Presumably the proof there could be adapted, with some additional complications arising from the fact that the contour is no longer the unit circle.

Since we were able to use a suitable limit as sum of residues to prove the involution property, this suggests looking at other such limits. The next simplest is the following, where up to symmetry we have residues at two points.

Proposition 13.38. *Let $g \in \mathcal{M}_{\eta,x_0,x_1,\dots,x_m} \subset \mathcal{M}_{\eta,x_0,x_0/u,x_1,\dots,x_m}$. Then*

$$\lim_{u\to 1/q} \int K(z,w;\eta,x_0,x_0/u)g(w) = (-1/pq^2x_0)\frac{z\theta_p(z/qx_0, z/q^2x_0, \eta/qzx_0, \eta/q^2zx_0)}{\theta_p(z^2/\eta)}(g(z)-g(qz)).$$

Here, we should recognize the right-hand side (apart from the factor $-1/pq^2x_0$) as the action of the standard operator of degree s. In general, taking the limit $u \to q^{-l}$ gives the operator of degree ls; we omit the details.

Now, suppose the parameters satisfy $\eta/x_ix_j \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$ for $1 \le i,j \le m$ and $x_i/x_j \notin p^{\mathbb{Z}}q^{\mathbb{Z}}$, $1 \le i < j \le m$.

Theorem 13.39. *With parameters as above, there is a subsheaf $\widetilde{\mathcal{M}er} \subset \mathcal{M}er$ such that*

$$\mathrm{Hom}(\mathcal{O}_X(-ds-d'f+r_1e_1+\cdots+r_me_m), \widetilde{\mathcal{M}er}) = \mathcal{M}_{q^{-d}\eta,q^{d-d'}x_0,q^{-r_1}x_1,\dots,q^{-r_m}x_m;p,q} \tag{13.111}$$

Proof. We first show that the given family of spaces forms a module over $\mathcal{S}'_{\eta,x_0,\dots,x_m;p,q}$. Equivalently, we need to show that if $D = ds + d'f - r_1e_1 - \cdots - r_me_m$, then

$$[D]\mathcal{M}_{\eta,x_0,x_1,\dots,x_m;p,q} \subset \mathcal{M}_{q^{-d}\eta,q^{d-d'}x_0,q^{-r_1}x_1,\dots,q^{-r_m}x_m;p,q}. \tag{13.112}$$

The case $D = e_1$ follows from

$$\mathcal{M}_{\eta,x_0,x_1,x_2,\dots,x_m;p,q} \subset \mathcal{M}_{\eta,x_0,qx_1,x_2,\dots,x_m;p,q}, \tag{13.113}$$

and the case $D = f - e_1$ is nearly as straightforward. The case $D = s$ is more complicated. To deal with this case, we must consider the function

$$h(z) = \frac{z\theta_p(z/qx_0, z/q^2x_0, \eta/qzx_0, \eta/q^2zx_0)}{\theta_p(z^2/\eta)}(g(z)-g(qz)). \tag{13.114}$$

This certainly satisfies the modified symmetry condition $h(p\eta/z) = h(z)$, and

$$\Gamma_{p,q}(z/pq^2x_0, \eta/q^2zx_0)\prod_{1\le i\le m}(z/x_i, p\eta/x_iz; p, q)h(z) \tag{13.115}$$

is clearly holomorphic outside points with $z^2 \in p^{\mathbb{Z}}\eta$. That it is holomorphic at $p^{i/2}\sqrt{\eta}$ then follows from the strange condition in the special case $g(p^{1/2}q\sqrt{\eta}) = g(p^{1/2}\sqrt{\eta})$. It thus remains to show that h satisfies the strange condition, in the form $h(p^{(1+i)/2}q^{(j)/2}\sqrt{\eta}) = h(p^{(1+i)/2}q^{(-j)/2}\sqrt{\eta})$. By the symmetry of h, this is equivalent to $h(p^{(1+i)/2}q^{(j)/2}\sqrt{\eta}) = h(p^{(1-i)/2}q^{j/2}\sqrt{\eta})$, which reduces easily to the corresponding conditions for the factors

$$\frac{z\theta_p(z/qx_0, z/q^2x_0, \eta/qzx_0, \eta/q^2zx_0)}{\theta_p(z^2/\eta)} \tag{13.116}$$

and $g(z) - g(qz)$.

Now, the conditions on the parameters allow us to apply any reflection in A_{m-1}, implying that operators of degree $[e_i]$ and $[f-e_i]$ also act correctly. This also ensures that $\mathcal{S}'_{\eta,x_0,x_1;q;p}$ is generated by $[s]$, $[e_1]$ and $[f-e_1]$, and thus that

$$[ds+d'f-r_1e_1]\mathcal{M}_{\eta,x_0,\dots,x_m;p,q} \subset \mathcal{M}_{q^{-d}\eta,q^{d-d'}x_0,q^{-r_1}x_1,x_2,\dots,x_m;p,q}, \tag{13.117}$$

and similarly for $[ds + d'f - r_i e_i]$. More generally,

$$[ds + d'f - r_1 e_1 - \cdots - r_m e_m]\mathcal{M}_{\eta,x_0,\ldots,x_m;p,q} \subset [ds + d'f - r_i e_i]\mathcal{M}_{\eta,x_0,\ldots,x_m;p,q} \subset \mathcal{M}_{q^{-d}\eta,q^{d-d'}x_0,\ldots,x_{i-1},q^{-r_i}x_i,x_{i+1},\ldots,x_m;p,q}. \quad (13.118)$$

Since

$$\bigcap_{1\le i\le m} \mathcal{M}_{q^{-d}\eta,q^{d-d'}x_0,\ldots,x_{i-1},q^{-r_i}x_i,x_{i+1},\ldots,x_m;p,q} = \mathcal{M}_{q^{-d}\eta,q^{d-d'}x_0,q^{-r_1}x_1,\ldots,q^{-r_m}x_m;p,q}, \quad (13.119)$$

it follows that this indeed gives a module as required.

It remains only to show that this module is saturated (it is a submodule of $\mathcal{M}er$, so does not contain torsion). The conditions on the parameters ensure that $C_m + mf$ is ample, and thus it suffices to show saturation relative to elements of degree $C_m + mf$. Since $\mathcal{M}er$ is saturated, this reduces to showing that if g is a $pq\eta$-symmetric meromorphic function such that

$$[C_m + mf]g \subset \mathcal{M}_{\eta/q^2,x_0/q^{m+1},x_1/q,x_2/q,\ldots,x_m/q;p,q}, \quad (13.120)$$

then $g \in \mathcal{M}_{\eta,x_0,x_1,x_2,\ldots,x_m;p,q}$. Since $T[mf] \subset [C_m + mf]$, we conclude that

$$[mf]g \subset T^{-1}\mathcal{M}_{\eta/q^2,x_0/q^{m+1},x_1/q,x_2/q,\ldots,x_m/q;p,q} = \mathcal{M}_{\eta,x_0/q^m,x_1,x_2,\ldots,x_m;p,q}, \quad (13.121)$$

and thus reduce to showing that if $[f]g \subset \mathcal{M}_{\eta,x_0/q,x_1,x_2,\ldots,x_m;p,q}$, then $g \in \mathcal{M}_{\eta,x_0,\ldots,x_m;p,q}$. This is straightforward: the symmetry condition and the strange condition are unchanged under multiplication by symmetric elliptic theta functions, so only holomorphy could be at issue. Dividing by a general section of $[f]$ at worst introduces a single $q^{\mathbb{Z}}$-orbit of poles, but these cannot actually arise, since no other section of $[f]$ would cancel them. □

This sheaf is essentially $W(E_{m+1})$-invariant. More precisely, for every simple reflection, the action of that simple reflection on $\mathcal{M}$ (where defined) induces a semilinear action on the corresponding sheaves, compatible with the corresponding isomorphism of surfaces. Again, it suffices to consider the action on $[e_1]$, $[f - e_1]$ and $[s]$, which is mostly straightforward. Indeed, in each case (including the cases involving the Fourier transform), we can obtain the requisite identity as a limit of a braid relation in a suitable family of actions of $W(E_{m+2})$, using Proposition 13.38 and analogous limits for other reflections.

Now, the definition of the degree 0 subspace of $\widetilde{\mathcal{M}er}$ is invariant under swapping p and q. It follows that as η, $x_0, \ldots, x_m$ vary over their $p^{\mathbb{Z}}q^{\mathbb{Z}}$ orbits, the corresponding spaces not only give us modules over $\mathcal{S}'_{\eta,x_0,\ldots,x_m;p,q}$ but also over $\mathcal{S}'_{\eta,x_0,\ldots,x_m;q,p}$. As a result, this construction gives rise to a quasicoherent sheaf on $X_{\eta,x_0,\ldots,x_m;p,q} \times X_{\eta,x_0,\ldots,x_m;q,p}$, which is precisely the sheaf M_{Sol} we wished to construct above.

Since this sheaf M_{Sol} respects the action of $W(E_{m+1})$ where things are defined, we obtain the following.

Proposition 13.40. *The elliptic Riemann-Hilbert correspondence commutes with the action of $W(D_m)$.*

Proposition 13.41. *Suppose that $\rho(s-e_1) \notin \langle p, q\rangle$ and both $X_{\rho;q;p}$ and $X_{s_{s-e_1}\rho;q,p}$ admit Riemann-Hilbert functors. Then $s_{s-e_1} \operatorname{Sol}_{\rho;q;p} \cong \operatorname{Sol}_{s_{s-e_1}\rho;q;p} s_{s-e_1}$.*

13.5 Generalized hypergeometric functions

If M is a sheaf such that $\chi(M) = \chi(\mathrm{Sol}(M)) = 1$, then we would typically expect that both M and $\mathrm{Sol}(M)$ have unique global sections, which would combine to give a specific solution of the corresponding straight-line equation (which would also be a solution of the straight-line p-difference equation). Of course, in general, we could only expect $\chi(M) = 1$ to imply $\dim\Gamma(M) = 1$ on an open set, but there is an important special case. Indeed, in the rigid case, if $M \cong \mathcal{O}_r$ where r is a root of $W(E_{m+1})$, then we automatically have $H^1(M) = 0$, and thus $\Gamma(M) \cong \mathbb{C}$. More precisely, if r is a root of $W(E_{m+1})$ such that $r \cdot f > 0$, then for sufficiently general ρ such that $pq\rho(r) = 1$, there is a sheaf $\mathcal{O}_r$ on $X_{\rho;q;p}$, and $\mathrm{Sol}(\mathcal{O}_r) \cong \mathcal{O}_r$. In particular, any such root gives rise to a natural family of functions (modulo scalars), and it is natural to wonder whether we can give an explicit description of this family.

In the simplest case $r = s - e_1$, this is quite simple: indeed, the corresponding straight-line equation is just $[s - e_1]g = 0$, and the function 1 clearly satisfies both the q-difference and the p-difference equation (and, more importantly, has the correct singularities).

More generally, if $G_r(z;\rho;p,q)$ is the solution corresponding to the root r and s is any admissible simple reflection, then it follows immediately from our general theory that

$$G_{sr}(z;\rho;p,q) \propto sG_r(z;s^{-1}\rho;p,q). \tag{13.122}$$

Since those simple reflections which do not act trivially either multiply by a product of Gamma functions or integrate against the standard kernel, we conclude the following.

Proposition 13.42. *For any root $r \in W(E_{m+1})$, there is a nonzero meromorphic function (a* generalized elliptic hypergeometric function*) $G_r(z;\rho;p,q) = G_r(z;\rho;q,p)$ on the subspace of parameters with $pq\rho(r) = 1$, expressible as a (possibly multiple) contour integral of a product of elliptic Γ functions, and such that $[r]G_r(z;\rho;p,q) = 0$.*

Proof. Note that although the Fourier transform was only defined as an integral when certain parameter conditions were satisfied, it follows from [49, Thm. 10.2] that the iterated integrals we obtain are meromorphic on the full parameter space. □

Remark. In other words, there is an integral representation for a solution to the equation $[r]g = 0$. Note, however, that we do not claim that this function is in any way unique; indeed, we will see that changing the path from $s - e_1$ to r can multiply G_r by a product of elliptic gamma functions involving the parameters (but not z).

Remark. For any element $w \in W(E_{m+1})$, $w \notin W(D_m)$, the corresponding action on meromorphic functions is given by an iterated integral transformation, which by Fubini can be expressed as a univariate integral transformation in which the kernel is itself expressed as an iterated integral. In fact, the resulting kernel is (proportional to) a special case of an appropriate generalized elliptic hypergeometric function!

The simplest cases of this are the first-order equations. Applying suitable reflections in $W(D_m)$ to $s - e_1$ gives us the following expressions (where we set $\eta = 1/pq$ for simplicity):

$$G_{s+f-e_1-e_2-e_3}(z;1/pq,x_0,x_1,x_2,x_3;p,q) = \frac{\Gamma_{p,q}(z^{\pm1}/x_1,z^{\pm1}/x_2,z^{\pm1}/x_3)}{\Gamma_{p,q}(z^{\pm1}/pqx_0)}, \quad x_0 = x_1x_2x_3$$

$$G_{s+2f-e_1-\cdots-e_5}(z;1/pq;x_0,x_1,\ldots,x_5;p,q) = \frac{\Gamma_{p,q}(z^{\pm1}/x_1,\ldots,z^{\pm1}/x_5)}{\Gamma_{p,q}(z^{\pm1}/pqx_0)}, \qquad x_0 = pqx_1\cdots x_5$$

$$G_{s+3f-e_1-\cdots-e_7}(z;1/pq;x_0,x_1,\ldots,x_7;p,q) = \frac{\Gamma_{p,q}(z^{\pm1}/x_1,\ldots,z^{\pm1}/x_7)}{\Gamma_{p,q}(z^{\pm1}/pqx_0)}, \qquad x_0 = p^2q^2x_1\cdots x_7$$

and so forth. Permuting so that $e_1 \cdot r = 0$ and then reflecting in $s - e_1$ gives a sequence of univariate integrals:

$$G_{s+f-e_2-e_3-e_4}(z; 1/pq, x_0, x_1, x_2, x_3, x_4; p, q)$$
$$= \frac{\Gamma_{p,q}(z^{\pm 1}/x_1)}{\Gamma_{p,q}(x_1/x_0, z^{\pm 1}/pqx_0)} I^{(0)}(pq\sqrt{x_0 x_1}, \sqrt{x_0/x_1}/x_2, \sqrt{x_0/x_1}/x_3, \sqrt{x_0/x_1}/x_4, \sqrt{x_1/x_0} z^{\pm 1}; p, q)$$

with $x_2 x_3 x_4 = x_0$,

$$G_{2s+2f-e_1-e_2-e_3-e_4-e_5-e_6}(z; 1/pq, x_0, x_1, x_2, x_3, x_4, x_5, x_6; p, q)$$
$$= \frac{\Gamma_{p,q}(z^{\pm 1}/x_1)}{\Gamma_{p,q}(x_1/x_0, z^{\pm 1}/pqx_0)} I^{(1)}(pq\sqrt{x_0 x_1}, \sqrt{x_0/x_1}/x_2, \ldots, \sqrt{x_0/x_1}/x_6, \sqrt{x_1/x_0} z^{\pm 1}; p, q)$$

with $pqx_1 x_2 \cdots x_6 = x_0^2$,

$$G_{3s+3f-2e_1-e_2-e_3-e_4-e_5-e_6-e_7-e_8}(z; 1/pq, x_0, x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8; p, q)$$
$$= \frac{\Gamma_{p,q}(z^{\pm 1}/x_1)}{\Gamma_{p,q}(x_1/x_0, z^{\pm 1}/pqx_0)} I^{(2)}(pq\sqrt{x_0 x_1}, \sqrt{x_0/x_1}/x_2, \ldots, \sqrt{x_0/x_1}/x_8, \sqrt{x_1/x_0} z^{\pm 1}; p, q)$$

with $p^2 q^2 x_1 x_2 \cdots x_8 = x_0^2$, and so forth, where $I^{(m)}$ denotes the "order m elliptic beta integral", see in particular [51], which derived the matrix form of the q-difference equations satisfied by $I^{(m)}$. In particular, the limits of $I^{(m)}$ include the ordinary hypergeometric function ${}_m F_{m-1}$, justifying the name "generalized elliptic hypergeometric function" above.

Note that for $m = 0$, we did not actually change the root when reflecting in $s - e_1$, and thus (for $x_2 x_3 x_4 = x_0$)

$$\frac{\Gamma_{p,q}(z^{\pm 1}/x_1)}{\Gamma_{p,q}(z^{\pm 1}/pqx_0)\Gamma_{p,q}(x_1/x_0)} I^{(0)}(pq\sqrt{x_0 x_1}, \sqrt{x_0/x_1}/x_2, \sqrt{x_0/x_1}/x_3, \sqrt{x_0/x_1}/x_4, \sqrt{x_1/x_0} z^{\pm 1}; p, q)$$
$$\propto \frac{\Gamma_{p,q}(z^{\pm 1}/x_2, z^{\pm 1}/x_3, z^{\pm 1}/x_4)}{\Gamma_{p,q}(z^{\pm 1}/pqx_0)}. \tag{13.123}$$

Using the explicit evaluation of the elliptic beta integral, we find that the constant is

$$\frac{\Gamma_{p,q}(x_2/x_1, x_3/x_1, x_4/x_1)}{\Gamma_{p,q}(x_2/x_0, x_3/x_0, x_4/x_0)}. \tag{13.124}$$

In particular, we see as mentioned above that the value of G_r can change if we change the sequence of simple reflections taking $s - e_1$ to r. This path dependence can be measured by a 1-cocycle over the stabilizer in $W(E_{m+1})$ of $s - e_1$, which is generated by $s + f - e_1 - e_2 - e_3$, $f - e_1 - e_2$, and $e_i - e_{i+1}$ for $2 \leq i < m$, only the first of which has any effect on $G_{s-e_1} = 1$. We could remove this ambiguity if we could extend this cocycle to $W(E_{m+1})$ (in particular if we could trivialize the cocycle), as we could then rescale the action of the simple reflections in such a way that the stabilizer of $s - e_1$ fixed 1.

Note that the nontrivial transformation of $I^{(1)}$ corresponds to reflection in the simple root $e_1 - e_2$, though the above calculations only determine the dependence on z of the appropriate constant. A cleaner proof of the transformation using the above ideas is to compare expressions for the action an appropriate element of $W(E_{m+1})$ coming from different reduced words. We omit the details.

Appendix A An integrable system on relations

The original approach to noncommutative geometry was via generators and relations, an approach we have mostly eschewed above. We did, however, make brief use of a presentation for F_0 (valid whenever $s-f$ is ineffective) in proving flatness in that case. The key property of this presentation was that it gave us a quadratic Gröbner basis for the $\mathbb{Z}^2$-algebra, suggesting the general question of when such a presentation satisfies the Gröbner basis property.

In particular, suppose we have a presentation of a $\mathbb{Z}^2$-algebra in which starting at each object we have four generators (two each of degrees $(0,1)$ and $(1,0)$) and six relations (one each of degree $(2,0)$ and $(0,2)$, and four of degree $(1,1)$). If we view the horizontal relation as an element of the tensor product of the two Hom spaces, then it must have rank 2, lest the $\mathbb{Z}^2$-algebra have zero divisors. But then we can identify the two Hom spaces in such a way that the relation becomes commutation, and similarly for the vertical relations.

Thus the general form of the presentation involves generators

$$\begin{aligned} x_1^{(d')}, x_2^{(d')} &\in \mathrm{Hom}((d,d'),(d+1,d')) = \mathrm{Hom}((0,d'),(1,d')) \\ y_1^{(d)}, y_2^{(d)} &\in \mathrm{Hom}((d,d'),(d,d'+1)) = \mathrm{Hom}((d,0),(d,1)), \end{aligned}$$

which satisfy relations

$$\begin{aligned} x_1^{(d')} x_2^{(d')} &= x_2^{(d')} x_1^{(d')} \\ y_1^{(d)} y_2^{(d)} &= y_2^{(d)} y_1^{(d)} \\ x_i^{(d'+1)} y_j^{(d)} &= \sum_{k,l} A_{ij}^{(d,d')kl} y_k^{(d+1)} x_l^{(d')}, \end{aligned}$$

where $A^{(d,d')}$ is an invertible matrix for every d, d'. As in the elliptic case, the Hom spaces in quotient $\mathbb{Z}^2$-algebra are clearly spanned by products of the form $y_2^{a_1} y_1^{a_2} x_2^{a_3} x_1^{a_4}$, so have dimension at most $\max(d+1,0)\max(d'+1,0)$. The question is thus when these spanning sets are bases.

Note that we can rescale $A^{(d,d')}$ by rescaling one of the bases, and we have enough such freedom to allow us to rescale all of the $A^{(d,d')}$ independently. We will thus view each $A^{(d,d')}$ as a point in the appropriate projective space.

By the theory of noncommutative Gröbner bases, the ordered monomials will form a basis iff they form a basis in degree $(2,1)$ and $(1,2)$. To express this relation, it will be convenient to modify the parametrization of the relations slightly. The key idea is that a 2-dimensional vector space admits an essentially unique symplectic form ϵ allowing us to identify the space with its dual. This allows us to replace $A^{(d,d')}$ by an element of the tensor product of four 2-dimensional vector spaces. The relations then take the form

$$\begin{aligned} \epsilon^{ij} x_i^{(d')} x_j^{(d')} &= 0 \\ \epsilon^{ij} y_i^{(d)} y_j^{(d)} &= 0 \\ x_i^{(d'+1)} y_j^{(d)} &= \epsilon^{km}\epsilon^{ln} A_{ijkl}^{(d,d')} y_m^{(d+1)} x_n^{(d')}, \end{aligned}$$

where ϵ is the symplectic form (normalized so that $\epsilon^{12}=1$, say) and we take the convention from physics that any index that appears as both an upper and a lower index should be summed over. The Gröbner basis property in degree $(2,1)$ then says that modulo the A relations, we have

$$\langle (\epsilon^{ij} x_i^{(d'+1)} x_j^{(d'+1)}) y_1^{(d)}, (\epsilon^{ij} x_i^{(d'+1)} x_j^{(d'+1)}) y_2^{(d)} \rangle = \langle y_1^{(d+2)} (\epsilon^{ij} x_i^{(d')} x_j^{(d')}), y_2^{(d+2)} (\epsilon^{ij} x_i^{(d')} x_j^{(d')}) \rangle. \quad \text{(A.1)}$$

This induces an isomorphism between $\mathrm{Hom}((d,d'),(d+1,d'))$ and $\mathrm{Hom}((d,d'+2),(d+1,d'+2))$, so let us for the moment change basis in the latter so that the condition becomes

$$(\epsilon^{ij}x_i^{(d'+1)}x_j^{(d'+1)})y_k^{(d)} = y_k^{(d+2)}(\epsilon^{ij}x_i^{(d')}x_j^{(d')}). \tag{A.2}$$

We have

$$\begin{aligned}(\epsilon^{ij}x_i^{(d'+1)}x_j^{(d'+1)})y_k^{(d)} &= \epsilon^{ij}\epsilon^{ln}\epsilon^{mo}A_{jklm}^{(d,d')}x_i^{(d'+1)}y_n^{(d+1)}x_o^{(d')} \\ &= \epsilon^{ij}\epsilon^{ln}\epsilon^{mo}\epsilon^{eg}\epsilon^{fh}A_{jklm}^{(d,d')}A_{inef}^{(d+1,d')}y_g^{(d+2)}x_h^{(d')}x_o^{(d')},\end{aligned} \tag{A.3}$$

and thus the Gröbner basis condition is that

$$\epsilon^{ij}\epsilon^{ln}\epsilon^{mo}\epsilon^{eg}\epsilon^{fh}A_{jklm}^{(d,d')}A_{inef}^{(d+1,d')} = \delta_k^g\epsilon^{ho} \tag{A.4}$$

(up to a change of basis in g), which can be restated as

$$\epsilon^{ie}\epsilon^{kf}A_{ijkl}^{(d,d')}A_{efgh}^{(d+1,d')} = -\epsilon_{jg}\epsilon_{lh}, \tag{A.5}$$

where ϵ_{ij} is the inverse tensor, $\epsilon^{ij}\epsilon_{jk} = \delta_k^i$. Similarly, the other Gröbner basis property becomes

$$\epsilon^{ih}\epsilon^{jf}A_{ijkl}^{(d,d')}A_{efgh}^{(d,d'+1)} = -\epsilon_{kg}\epsilon_{le}. \tag{A.6}$$

Each of these can be easily rewritten as stating that two matrices are inverse to each other, and thus we see that $A^{(d,d')}$ uniquely determines $A^{(d\pm1,d')}$ and $A^{(d,d'\pm1)}$, up to the aforementioned choice of basis. We in fact find that either way of computing $A^{(d+1,d'+1)}$ gives the same answer, namely

$$A_{ijkl}^{(d+1,d'+1)} = A_{lkji}^{(d,d')}. \tag{A.7}$$

(This reduces easily to the fact that $AB = 1$ iff $BA = 1$.) It follows that the choices of basis made above are consistent as we move around the category. In particular, we obtain a unique $\mathbb{Z}^2$-algebra associated to any sufficiently general choice of $A^{(0,0)}$, in which all of the other relations are determined as rational functions of the original relation.

It turns out that the resulting $\mathbb{Z}^2$-algebra is generically just our elliptic F_0. Indeed (following [8]), if we choose a genus 1 curve C, two degree 2 line bundles $\mathcal{L}_1$, $\mathcal{L}_2$, a degree 0 line bundle q, and bases of $\Gamma(C;\mathcal{L}_1)$, $\Gamma(C;\mathcal{L}_2)$, $\Gamma(C;\mathcal{L}_1\otimes q)$, $\Gamma(C;\mathcal{L}_2\otimes q)$, then the isomorphism

$$\Gamma(C;\mathcal{L}_1)\otimes\Gamma(C;\mathcal{L}_2\otimes q) \cong \Gamma(C;\mathcal{L}_1\otimes(\mathcal{L}_2\otimes q)) \cong \Gamma(C;(\mathcal{L}_1\otimes q)\otimes\mathcal{L}_2) \cong \Gamma(C;\mathcal{L}_1\otimes q)\otimes\Gamma(C;\mathcal{L}_2) \tag{A.8}$$

induces a matrix of the form $A^{(0,0)}$, and one can moreover reconstruct the data from $A^{(0,0)}$. Indeed, C is reconstructed as the decomposable tensors (mod scalars) that map to decomposable tensors, and in particular inherits four maps to $\mathbb{P}^1$. The tuple $(C,\mathcal{L}_1,\mathcal{L}_2,q)$ varies over a 3-dimensional moduli space, and each choice of basis modulo scalars introduces 3 more parameters, for a total of 15. That is, the moduli space of tuples of data as above is 15-dimensional, and the map to $A^{(0,0)}$ embeds it in $\mathbb{P}^{15}$ as an open subset. The operation taking $A^{(0,0)}$ to $A^{(0,1)}$ is then just the operation of multiplying $\mathcal{L}_1$ by q, and the fact that the new space has an induced basis arises from the natural isomorphism

$$\Gamma(C;\mathcal{L}_1\otimes q^2) \cong \Gamma(C;\mathcal{L}_1) \tag{A.9}$$

coming from translation by q.

Note that this gives a possible approach to understanding degenerations of the noncommutative $\mathbb{P}^1\times\mathbb{P}^1$: simply understand the natural stratification of the complement of the elliptic locus in $\mathbb{P}^{15}$.

This is actually quite doable, but since it only gives cases admitting Fourier transformations, this is not the approach we take in [44]. We do, however, mention that the case in which $A^{(0,0)}$ corresponds to the relations

$$\begin{aligned} x_1y_1 &= y_1x_1 \\ x_1y_2 &= y_2x_1 \\ x_2y_1 &= y_1x_2 \\ x_2y_2 &= y_2x_2 + y_1x_1 \end{aligned} \tag{A.10}$$

is a fixed point for the evolution, and thus the resulting $\mathbb{Z}^2$-algebra comes from a bigraded algebra. This is essentially a compactification of the Weyl algebra $Dt - tD = 1$, which we can see by taking $(x_1, x_2) = (1, D)$, $(y_1, y_2) = (1, t)$. The analogue of the generalized Fourier transform gives relations

$$\begin{aligned} x_1y_1 &= y_1x_1 \\ x_1y_2 &= y_2x_1 \\ x_2y_1 &= y_1x_2 \\ x_2y_2 &= y_2x_2 - y_1x_1, \end{aligned} \tag{A.11}$$

but of course these are equivalent to the original relations by taking $y_2 \mapsto -y_2$. In other words, the generalized Fourier transformation acts on this algebra as $(D, t) \mapsto (-t, D)$, justifying the name.

We now wish to consider the matrix-valued rational functions $A^{(d,d')} \in \mathrm{GL}_4(k(A^{(0,0)}))$ in more detail. In light of equation (A.7), this essentially reduces to the case $d = 0$. Here, we may note that the expression for $A^{(0,d')}$ in terms of $A^{(0,0)}$ is the d'-th iterate of the function giving $A^{(0,1)}$ in terms of $A^{(0,0)}$. The relation to arithmetic on a family of abelian surfaces $\mathrm{Pic}^0(C)^2$ suggests that this should be integrable (in the algebraic entropy sense, see below), in that the degrees of these rational functions should grow unexpectedly slowly. Not only is this true, but we can in fact be quite precise about the degrees.

Proposition A.1. *Let T be the birational map $\mathbb{P}^{15} \to \mathbb{P}^{15}$ such that $A^{(0,1)} \propto TA^{(0,0)}$. Then T^d has degree $2d^2 + 1$.*

To prove this, we note that T fits into a larger family of birational maps coming from the above structure. Indeed, T can be expressed in terms of the obvious action of S_4 together with the birational involution R_{12} defined by

$$\epsilon^{ie}\epsilon^{jf}A_{ijkl}(R_{12}A)_{efgh} = -\epsilon_{kg}\epsilon_{lh}. \tag{A.12}$$

The following is again straightforward.

Lemma A.2. *The map R_{12} commutes with the subgroup $\langle (12), (34), (14)(23) \rangle$ of S_4.*

We thus see that R_{12} has three distinct conjugates under S_4, which we may denote by $R_1 = R_{23}$, $R_2 = R_{13}$ and $R_3 = R_{12}$. We thus obtain a representation of $W_3 \rtimes S_4$ as birational automorphisms of $\mathbb{P}^{15}$, where W_3 denotes the free Coxeter group on three generators (acting as R_1, R_2, R_3. Since any element of S_4 simply permutes the coordinates, to understand degrees of elements in this group, it suffices to understand degrees of elements in W_3. Note that as a Coxeter group, W_3 comes with a reflection representation ρ, in which the generators act as

$$s_1 = \begin{pmatrix} -1 & 2 & 2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad s_2 = \begin{pmatrix} 1 & 0 & 0 \\ 2 & -1 & 2 \\ 0 & 0 & 1 \end{pmatrix}, \quad s_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 2 & -1 \end{pmatrix}, \tag{A.13}$$

i.e., as reflections in the coordinate vectors relative to the inner product

$$\begin{pmatrix} 2 & -2 & -2 \\ -2 & 2 & -2 \\ -2 & -2 & 2 \end{pmatrix}. \tag{A.14}$$

We then have the following.

Theorem A.3. *For any element $g \in W_3$, let T_g denote the corresponding birational automorphism of $\mathbb{P}^{15}$. Then the degree of T_g is 2 less than the sum of the coefficients of $\rho(g)$.*

Proof. For each i, R_i is essentially just taking the inverse transpose of an appropriate matrix (relative to the symmetric form $\epsilon \otimes \epsilon$). Let $\delta_i(A)$ denote the determinant of that matrix, and let Σ_i denote the (smooth!) subscheme $\mathbb{P}^3 \times \mathbb{P}^3 \subset \mathbb{P}^{15}$ where that matrix has rank 1.

Now, for $d, r_1, r_2, r_3 \in \mathbb{N}$, let $V(d, r_1, r_2, r_3)$ denote the space of homogeneous polynomials of degree d that vanish to order at least r_i along Σ_i. We first claim that if $\tilde{R}_1$ denotes the primitive polynomial representation of R_1, one has

$$\tilde{R}_1 V(d, r_1, r_2, r_3) \subset \delta_1(A)^{r_1} V(3d - 4r_1, 2d - 3r_1, r_2, r_3). \tag{A.15}$$

Indeed, inverse transpose respects tensor products, so that $\tilde{R}_1$ is defined on and preserves Σ_2 and Σ_3, and therefore maps the corresponding ideals into themselves. Similarly, since the coefficients of $\tilde{R}_1$ are (up to sign and ordering) just the 3×3 minors of the relevant matrix, they have degree 3 and vanish to order 2 along Σ_1. It follows that

$$\tilde{R}_1 V(d, r_1, r_2, r_3) \subset \tilde{R}_1 V(d, 0, r_2, r_3) \subset V(3d, 2d, r_2, r_3). \tag{A.16}$$

Now, the ideal of Σ_1 is generated by 2×2 minors of the appropriate matrix, and the image of such a minor under $\tilde{R}_i$ is the product of the complementary minor by the determinant. We thus find that any element of $\tilde{R}_1 V(d, r_1, r_2, r_3)$ is a multiple of $\delta_1(A)^{r_1}$. Since this has degree 4 and vanishes to order 3 along Σ_1, we conclude

$$\tilde{R}_1 V(d, r_1, r_2, r_3) \subset \delta_1(A)^{r_1} V(3d - 4r_1, 2d - 3r_1, r_2, r_3) \tag{A.17}$$

as required.

Since $\tilde{R}_1^2(f) = \delta_1^{2\deg(f)} f$, $\tilde{R}_1 \delta_1(A) = \delta_1(A)^3$, we further have

$$\begin{aligned} \delta_1(A)^{2d} V(d, r_1, r_2, r_3) &= \tilde{R}_1^2 V(d, r_1, r_2, r_3) \\ &\subset \tilde{R}_1 \delta_1(A)^{r_1} V(3d - 4r_1, 2d - 3r_1, r_2, r_3) \\ &\subset \delta_1(A)^{3r_1} \delta_1(A)^{2d - 3r_1} V(d, r_1, r_2, r_3) \\ &= \delta_1(A)^{2d} V(d, r_1, r_2, r_3) \end{aligned} \tag{A.18}$$

Both inclusions are therefore equalities, and thus since the map on parameters is an involution, we conclude that

$$\tilde{R}_1 V(d, r_1, r_2, r_3) = \delta_1(A)^{r_1} V(3d - 4r_1, 2d - 3r_1, r_2, r_3), \tag{A.19}$$

and similarly for $\tilde{R}_2$, $\tilde{R}_3$. Note also that if the elements of $V(d, r_1, r_2, r_3)$ are relatively prime, then the same will be true for $V(3d - 4r_1, 2d - 3r_1, r_2, r_3)$, as otherwise we could apply $\tilde{R}_1$ to transport the gcd back.

Note that if we define $W(r_1, r_2, r_3) := V(r_1 + r_2 + r_3 - 2, r_1 - 1, r_2 - 1, r_3 - 1)$ for $r_1, r_2, r_3 \geq 1$, then the above action becomes

$$\tilde{R}_1 W(r_1, r_2, r_3) = \delta_1(A)^{r_1 - 1} W(-r_1 + 2r_2 + 2r_3, r_2, r_3). \tag{A.20}$$

Here, of course, the action on the parameters is precisely the reflection representation.

Now, if $s_{i_1} s_{i_2} \cdots s_{i_m}$ is any word in W_3, the corresponding birational map is $\tilde{R}_{i_1} \cdots \tilde{R}_{i_m}$, and the desired degree is the residual degree after removing the common factors from

$$\tilde{R}_{i_1} \cdots \tilde{R}_{i_m} V(1, 0, 0, 0) = \tilde{R}_{i_1} \cdots \tilde{R}_{i_m} W(1, 1, 1) \tag{A.21}$$

After removing the gcd, we are left with a space of the form $W(r_1, r_2, r_3)$ where (r_1, r_2, r_3) is the image of $(1, 1, 1)$ under $\rho(s_{i_1} \cdots s_{i_m})$, and the result thus follows from the fact that elements of $W(r_1, r_2, r_3)$ have degree $r_1 + r_2 + r_3 - 2$. □

Remark. Relative to the above identification of an open subset of $\mathbb{P}^{15}$ with a moduli space of smooth genus 1 curves with additional data, we see that W_3 acts as automorphisms on the corresponding abelian surface $\mathrm{Pic}^0(C)^2$. It follows that the generic W_3 orbit is dense in a Kummer surface, namely $\mathrm{Pic}^0(C)^2/\pm 1$. One can verify that each such Kummer surface is a quartic hypersurface in some $\mathbb{P}^3 \subset \mathbb{P}^{15}$ with 12 singular points. The above argument arose from considerations of how W_3 acts on monodromy-invariant divisors on such Kummers; in particular, the union over all such Kummers of the singular points turns out to have closure precisely $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3$.

To prove Proposition A.1, we need simply compute $\rho((s_1 s_2)^k)$ and $\rho((s_1 s_2)^k s_1)$, as the corresponding S_4 double cosets contain T^{2k} and T^{2k+1} respectively. The claim thus reduces to showing

$$\rho((s_1 s_2)^k) = \begin{pmatrix} 2k+1 & -2k & 4k^2 + 2k \\ 2k & 1 - 2k & 4k^2 - 2k \\ 0 & 0 & 1 \end{pmatrix}, \tag{A.22}$$

easily verified by induction.

Recall that the *algebraic entropy* of a birational automorphism T of projective space is given by [7]

$$\lim_{n \to \infty} \frac{1}{n} \log \deg(T^n), \tag{A.23}$$

with an automorphism considered integrable if its algebraic entropy vanishes. The above result makes it straightforward to compute the algebraic entropy of any birational automorphism in the above group. This is most simply expressed in terms of the homomorphism $\phi : W_3 \rtimes S_4 \to \mathrm{PGL}_2(\mathbb{Z})$ induced by the classical isomorphism $W_3 \rtimes S_3 \cong \mathrm{PGL}_2(\mathbb{Z})$. Although elements of $\mathrm{PGL}_2(\mathbb{Z})$ do not have eigenvalues per se, their preimages in $\mathrm{GL}_2(\mathbb{Z})$ have eigenvalues, and the choice of preimage only affects the sign; as a result, we may define *absolute values* of eigenvalues of elements in $\mathrm{PGL}_2(\mathbb{Z})$.

Corollary A.4. *For any $g \in W_3 \rtimes S_4$, let λ denote the largest absolute value of an eigenvalue of $\phi(g)$. Then T_g has algebraic entropy $2 \log \lambda$.*

Proof. We first observe that $T_g = \frac{1}{12} T_{g^{12}}$; since $g^{12} \in W_3$ for any $g \in W_3 \rtimes S_4$, we thus reduce to the case $g \in W_3$. Moreover, the reflection representation of W_3 is just (up to rational change of basis) the restriction of the adjoint representation of $\mathrm{PGL}_2(\mathbb{Z})$ (i.e., the representation corresponding to the symmetric square of the 2-dimensional representation of $\mathrm{GL}_2(\mathbb{Z})$). If some preimage of $\phi(g)$ has eigenvalues $z, \pm 1/z$, then the symmetric square has eigenvalues $z^2, \pm 1, 1/z^2$. It follows that the largest eigenvalue of $\rho(g)$ has absolute value λ^2. It follows that the matrix coefficients of $\rho(g^n)$ are $O(\lambda^{2n} n^2)$, and thus the same holds for $\deg(T_g^n)$. The claim for the algebraic entropy follows immediately. □

If we attempt to give a similar presentation for a deformation of $(\mathbb{P}^1)^3$, then in addition to each plane having to satisfy the above Gröbner basis property, there is also an additional consistency condition in degree $(1,1,1)$, comparing the two ways of reducing $z_i y_j x_k$. A calculation along the above lines gives a Yang-Baxter-type identity, of the form

$$(A \otimes 1_2)(1_2 \otimes B)(C \otimes 1_2) = (1_2 \otimes D)(E \otimes 1_2)(1_2 \otimes F) \tag{A.24}$$

where each of the six 4×4 matrices comes from the relations on a given face of the cube. If one moves to an adjacent cube, one of the relations is directly carried over, while four others are determined by the planar Gröbner basis property. This overdetermines the remaining matrix, which in fact depends only on the relation on the opposite face. In particular, the 2×2 elliptic solution of the dynamic Yang-Baxter equation [20] gives rise to a deformation of $(\mathbb{P}^1)^3$ generalizing the elliptic $\mathbb{P}^1 \times \mathbb{P}^1$, which further extends to $(\mathbb{P}^1)^n$. We omit the details, as we will give an equivalent construction below. We do mention, however, that the set of solutions of the above YBE is invariant under the transformation $(A, B, C, D, E, F) \mapsto (A^{-1}, D, E, B, C, F^{-1})$ as well as under linear transformations arising by rewriting the relation as

$$\epsilon^{gj}\epsilon^{hk}\epsilon^{il} A_{abgl} B_{cdhj} C_{efik} = \epsilon^{mp}\epsilon^{nq}\epsilon^{or} D_{bcmr} E_{fanp} F_{deoq} \tag{A.25}$$

and observing that this admits a permutation action of the dihedral group of order 12. In particular, there is an integrable action on $(\mathbb{P}^{15})^6$ (coming from a subgroup of the wreath product $W_3^6 \rtimes S_6$) that preserves the locus of solutions of the generalized YBE.

Appendix B Blowing up $\mathbb{P}^n$

In addition to $\mathbb{P}^1 \times \mathbb{P}^1$ and F_1, there is another case in which we have (at least generically) a particularly nice presentation of our category with a well-behave Gröbner basis. This is suggested by the following observation in the commutative setting.

Proposition B.1. *Let X_2 be the blowup of (commutative) $\mathbb{P}^2$ in the three coordinate vectors. Then the* $\mathrm{Pic}(X_2)$*-graded coordinate ring of X_2 is the free polynomial ring in six generators of degree e_1, e_2, e_3, $h - e_2 - e_3$, $h - e_1 - e_3$, and $h - e_1 - e_2$.*

Proof. Consider any divisor class $dh - r_1 e_1 - r_2 e_2 - r_3 e_3$. Elements of that degree are homogeneous polynomials of degree d that vanish to order at least r_1, r_2, r_3 at the three coordinate vectors. This space is clearly spanned by monomials, and $x^{a_1} y^{a_2} z^{a_3}$ is in the space iff $a_2 + a_3 \geq r_1$, $a_1 + a_3 \geq r_2$, $a_1 + a_2 \geq r_3$, corresponding to the decomposition

$$\begin{aligned} dh - r_1 e_1 - r_2 e_2 - r_3 e_3 = {} & (a_2 + a_3 - r_1) e_1 + (a_1 + a_3 - r_2) e_2 + (a_1 + a_2 - r_2) e_3 \\ & + a_1(h - e_2 - e_3) + a_2(h - e_1 - e_3) + a_3(h - e_1 - e_2). \end{aligned} \tag{B.1}$$

In particular, the dimension of the given homogeneous subspace of the multigraded coordinate ring is precisely equal to the number of monomials in generators of the given degrees. □

Remark. The reduction of dimension computations to combinatorial calculations reflects the fact that this particular rational surface is a toric variety.

Remark. Similarly, one finds that if one blows up two distinct points of $\mathbb{P}^2$, the resulting $\mathrm{Pic}(X_1)$-graded coordinate ring is the free polynomial ring in five generators, of degrees e_1, e_2, $h - e_1 - e_2$, $h - e_1$, and $h - e_2$. Note, however, that the generators of degree $h - e_1$ and $h - e_2$ are not unique (even when considered mod scalars), making the resulting deformed presentation more complicated.

In the noncommutative setting, we have the following.

Lemma B.2. *If the surface $X_{\eta,x_0,x_1,x_2;q;C}$ has no effective -2-curves, then $\mathcal{S}'_{\eta,x_0,x_1,x_2;q;C}$ is generated by morphisms of degree s, e_1, e_2, $s+f-e_1-e_2$, $f-e_2$, $f-e_1$. This gives a presentation for $\mathcal{S}'_{\eta,x_0,x_1,x_2;q;C}$ with 6 generators from each object, with 15 relations forming a quadratic Gröbner basis.*

Proof. We first need to show that for any divisor class D, we can span $[D]$ by products of the given degrees. If $D \cdot e_2 < 0$, then $[D] = [e_2][D-e_2]$, so we can proceed by induction in $D \cdot C_2$. By symmetry, the same applies to any of the six -1-classes. We thus reduce to the case that D is nef, and using the $W(E_3)$-symmetry, that D is universally nef. By our surjectivity results, this reduces to the cases $D \in \{f, s+f, s+2f-e_1, 2s+3f-e_1-e_2\}$. For the first three cases, it suffices to consider leading coefficients, and we have:

$$\begin{aligned}[f] &= [e_1][f-e_1] + [e_2][f-e_2] \\ [s+f] &= [s][f] + [s+f-e_1-e_2][e_1][e_2] \\ [s+2f-e_1] &= [f-e_1][s+f] + [e_2][s+f-e_1-e_2][f]\end{aligned}$$

Similarly, we certainly obtain all leading terms for $[2s+3f-e_1-e_2]$, and $T \in [s+2f-e_1][s+f-e_2]$ follows by taking suitable choices of v, v', w in the proof of Lemma 3.3.

Each of the spaces of degree a -1-class is 1-dimensional, and thus we have the 6 generators as required. The degree of each generator has intersection 1 with C_2, giving us a positive grading on the category. We find by flatness that the Hom spaces from any given object have Hilbert series $1/(1-t)^6$ relative to this grading, and thus there must be 15 quadratic relations.

In fact, the degree 2 space is spanned by ordered products of the form

$$[s]^{a_1}[e_1]^{a_2}[e_2]^{a_3}[s+f-e_1-e_2]^{b_1}[f-e_2]^{b_2}[f-e_1]^{b_3} \tag{B.2}$$

with $a_1+a_2+a_3+b_1+b_2+b_3 = 2$. Indeed, if e, e' are orthogonal -1-classes, then

$$[e+e'] = [e][e'] = [e'][e], \tag{B.3}$$

letting us put those classes into preferred order. There are three more possible sums of distinct -1-classes, for which we have

$$\begin{aligned}[s+f-e_1] &= [e_2][s+f-e_1-e_2] + [s][f-e_1] \\ [s+f-e_2] &= [e_1][s+f-e_1-e_2] + [s][f-e_2] \\ [f] &= [e_1][f-e_1] + [e_2][f-e_2],\end{aligned}$$

from which the claim follows.

By our Hilbert series calculation, these ordered products must in fact form a basis, and thus not only can there be no additional relations, but the quadratic relations must form a Gröbner basis. □

Remark. Note that the relations have a very similar form to the $\mathbb{P}^1 \times \mathbb{P}^1$ case, in that the relations are precisely those quadratic relations obtained from multiplication on C.

It turns out that we can use this description to define deformations of the blowup of $\mathbb{P}^n$ in the coordinate vectors for any $n \geq 2$. The argument of Proposition B.1 works *mutatis mutandis* to

show that the resulting multigraded algebra is freely generated by commuting elements of degree $e_1,\dots,e_{n+1}$, $f_1,\dots,f_{n+1}$, where we denote

$$f_i := e_i + (h - e_1 - \cdots - e_{n+1}). \tag{B.4}$$

We thus need to construct deformed relations that still form a Gröbner basis.

Let C be a smooth genus 1 curve, and let $x_1,\dots,x_n$, H, q be divisors on C of degree $1,\dots,1$, $n+1$, and 0 respectively; assume furthermore that no divisor of the form $q^l H/x_1\cdots x_{n+1}$ or $q^l x_i/x_j$ is principal. Also define a pairing on $\Lambda := \langle h, e_1, \cdots, e_{n+1}\rangle$ by

$$h^2 = n-1, \quad h\cdot e_i = 0, \quad e_i\cdot e_j = -\delta_{ij}. \tag{B.5}$$

Finally, define a linear functional κ on Λ by $\kappa(h) = n+1$ and $\kappa(e_i) = 1$, and note that $\kappa(f_i) = 1$ and $\kappa(v) \in v^2 + 2\mathbb{Z}$ for all $v \in \Lambda$.

Now, given $v = dh - r_1e_1 - \cdots - r_{n+1}e_{n+1} \in \Lambda$, let D_v denote the divisor

$$q^{\frac{-v^2-\kappa(v)}{2}} H^d \prod_{1\le i\le n+1} x_i^{-r_i} \tag{B.6}$$

of degree $\kappa(v)$. Then we may define a Λ-algebra $\mathcal{B}_{H,x_1,\dots,x_{n+1};q;C}$ such that

$$\mathcal{B}(v,w) = \mathcal{L}(D_w/D_v), \tag{B.7}$$

with composition given by multiplication of functions on C. Note that for any function f on xC, we can add $\operatorname{div}(f)$ to any of the defining divisors and obtain an isomorphic category. In particular, the isomorphism class of the category depends only on the images of the defining divisors in the Picard group.

Note that in addition to the obvious S_{n+1} symmetry, we also have an isomorphism

$$\mathcal{B}_{H,x_1,\dots,x_{n+1};q;C} \cong \mathcal{B}_{(n-1)H(x_1\cdots x_{n+1})^{2-n},H/x_2\cdots x_{n+1},\dots,H/x_1\cdots x_n;q;C} \tag{B.8}$$

acting on Λ via the reflection

$$\begin{aligned} e_i &\mapsto h - e_1 - \cdots - e_{n+1} + e_i \\ h &\mapsto nh - (n-1)e_1 - \cdots - (n-1)e_{n+1}. \end{aligned}$$

In addition, the sub-$\mathbb{Z}^{n+1}$-algebra with objects in $\langle h - e_{n+1}, e_1, \dots, e_n\rangle$ is isomorphic to

$$\mathcal{B}_{H/x_{n+1},x_1,\dots,x_n;q;C}. \tag{B.9}$$

Now, let $\mathcal{B}'_{H,x_1,\dots,x_{n+1};q;C}$ be the subcategory of $\mathcal{B}_{H,x_1,\dots,x_{n+1};q;C}$ generated by elements of degree $e_1,\dots,e_{n+1}$, $f_1,\dots,f_{n+1}$, all of κ-degree 1. (Here we define the *κ-degree* of a morphism of degree v to be $\kappa(v)$.) This collection of elements is invariant under the S_{n+1} symmetry as well as the additional reflection, and thus $\mathcal{B}'$ inherits these symmetries, as well as the subalgebra property.

Finally, let $\mathcal{A}_{H,x_1,\dots,x_{n+1};q;C}$ be the category obtained from $\mathcal{B}'_{H,x_1,\dots,x_{n+1};q;C}$ by removing all relations of κ-degree $\neq 2$; again, this inherits an $A_1 \times A_n$ symmetry and reduction to subalgebras.

Theorem B.3. *For any $v, w \in \Lambda$, we have the direct sum decomposition*

$$\mathcal{A}(v,w) = \bigoplus_{\substack{\alpha,\beta\in\mathbb{N}^{n+1} \\ \sum_i(\alpha_i e_i + \beta_i f_i) = w - v}} [e_1]^{\alpha_1}\cdots[e_{n+1}]^{\alpha_{n+1}}[f_1]^{\beta_1}\cdots[f_{n+1}]^{\beta_{n+1}}. \tag{B.10}$$

In particular, $\mathcal{A}_{H,x_1,\dots,x_{n+1};q;C}$ is a flat family of Λ-algebras.

Proof. Note first that one has an isomorphism

$$\mathcal{S}'_{\eta,x_0,x_1,x_2;q;C} \cong \mathcal{A}_{\eta x_0,x_0,x_1,x_2;q;C}; \tag{B.11}$$

indeed, one finds that $\mathfrak{D}_{\eta x_0,x_0,x_1,x_2;q;C}(v) - D_v$ is principal for all v, giving an identification on leading coefficients; since we have shown that $\mathcal{S}'$ has precisely the given quadratic relations, the claim follows. In particular, it follows that the Theorem holds in the case $n = 2$.

Now, the category $\mathcal{A}$ has $2n + 2$ generators from each object. If we map the generators to $\{1, \dots, n + 1\}$ by taking both e_i and f_i to i, then any d of the generators determine a subset of $\{1, \dots, n + 1\}$ of size at most d. We may then use the $A_1 \times A_n$ symmetry to make that subset contained in $\{1, \dots, d\}$, and then use the subalgebra property to let us work in any $n \geq d - 1$.

In particular, to understand the relations, we may work in any $n \geq 1$; since we know the result holds for $n = 2$, it follows that the given elements span.

It remains to show the Gröbner basis property, but for this it suffices to consider products of three generators. Again, it suffices to verify the Gröbner basis property for $n = 2$, where we already know it holds. □

Remark. The form of divisors D_v was of course chosen to match up with $\mathcal{S}'$, but one can show in general that for $n \geq 2$, D_v must have the given form in order for the category to have enough relations. We omit the details.

Remark. It is, in fact, not difficult to directly verify the Gröbner basis property for $n = 2$ (and thus avoid using any results from the main part of the paper), as this simply involves verifying in each case that the resulting elements remain linearly independent in $\mathcal{B}$.

Let Z_n denote the Λ-graded algebra which in degree $dh - r_1e_1 - \cdots - r_{n+1}e_{n+1}$ consists of homogeneous polynomials of degree d on $\mathbb{P}^n$ vanishing at the ith coordinate vector with multiplicity at least r_i. This becomes a Λ-algebra $\mathcal{Z}_n$ in the usual way.

Proposition B.4. *There is an isomorphism $\mathcal{A}_{H,x_1,\dots,x_{n+1};1;C} \cong \mathcal{Z}_n$ acting trivially on objects.*

Proof. We may as well assume that H, $x_1,\dots,x_{n+1}$ are all effective. We can then use H to embed C in $\mathbb{P}^n$ in such a way that the images of $x_1,\dots,x_{n+1}$ are the coordinate vectors. We in particular find that there is a natural identification between the κ-degree 1 elements of the two categories; since the relations in $\mathcal{A}$ simply state that multiplication commutes, this identification extends to a functor. But we can argue as in Proposition B.1 to see that Z_n has no relations other than commutativity, and thus has the same relations as $\mathcal{A}$. □

The following is an immediate consequence.

Corollary B.5. *The subcategory of $\mathcal{A}_{H,x_1,\dots,x_{n+1};q;C}$ with objects $\mathbb{Z}h$ is a flat deformation of the $\mathbb{Z}$-algebra form of the homogeneous coordinate ring of $\mathbb{P}^n$.*

It turns out that this is not a new deformation of $\mathbb{P}^n$, as it can be seen to be a $\mathbb{Z}$-algebra version of the algebra $Q_n(\mathcal{E}, \eta)$ constructed by Feigin and Odesskii, see [38]. In fact, one finds that the above basis of $\mathcal{A}(0, dh)$ identifies that module with a module over $Q_n(\mathcal{E}, \eta)$ of the form considered in [38, Prop. 8]. An interesting consequence is that the resulting $\mathbb{Z}$-algebra is independent of $x_1,\dots,x_{n+1}$, suggesting that it may be possible to extend the definition of $\mathcal{A}$ to more general point configurations by taking a flat limit as in Section 7.

Though the deformation is not new, the extra scaffolding allows us to prove some embedding results. Note that the first two facts generalize the fact that F_1 is a blowup of $\mathbb{P}^2$ and the fact that $\mathbb{P}^1 \times \mathbb{P}^1$ is a quotient of the Sklyanin algebra.

Proposition B.6. *For $d \geq 1$, the subalgebra of $\mathcal{S}'_{\eta,x_0;q;C}$ with objects in $\mathbb{Z}(s+df)$ is a quotient of the above noncommutative $\mathbb{P}^{2d}$.*

Proposition B.7. *For $d \geq 1$, the subalgebra of $\mathcal{S}_{\eta,\eta';q;C}$ with objects in $\mathbb{Z}(s+df)$ is a quotient of the above noncommutative $\mathbb{P}^{2d+1}$.*

Proposition B.8. *The second Veronese of the elliptic noncommutative $\mathbb{P}^2$ is a quotient of the above noncommutative $\mathbb{P}^5$.*

Proof. In each case, let D be the given divisor, and consider the surface obtained by blowing up $n+1 = -D \cdot K$ points in general position; in particular, we insist that no two of the points are in the same $q^{\mathbb{Z}}$-orbit, and that $D - e_1 - \cdots - e_{n+1}$ not be a -2-curve. Then there is a map from $\langle h, e_1, \ldots, e_{n+1}\rangle$ to $\mathrm{Pic}(X_{n+1})$ given by $h \mapsto D$ and $e_i \mapsto e_i$, which moreover preserves the pairing and takes κ to intersection with $-K$. In particular, we find that the leading coefficient map is injective on Hom spaces in $\mathcal{S}$ or $\mathcal{S}'$ of degree corresponding to sums of one or two of the e_i or f_i. Since the above map on objects extends to a functor on the saturated leading coefficient categories and the relations come from leading coefficients, it follows that restricting to $\mathbb{Z}h$ gives a functor as required.

To see that this functor is surjective, we note that the target category is in each case generated in degree 1, and thus it suffices to prove surjectivity in degree h. But in each case the leading coefficient map is an isomorphism in the appropriate degree, and thus the result follows. □

Remark. The above argument works in any case such that the leading coefficient map identifies $\Gamma(\mathcal{O}_X(D))$ with the global sections of a very ample line bundle on C. However, for such a D, $D+K$ is ineffective, while D is nef, and consideration of the universally nef case shows that D must be equivalent under $W(E_{m+1})$ to one of the above three forms. In particular, we find that these are the only cases in which D is ample.

Remark. Note that although the noncommutative $\mathbb{P}^n$ is the $\mathbb{Z}$-algebra associated to a graded algebra, this is not true for the quotients in the Hirzebruch surface cases; the ideal is not invariant under the relevant automorphism of the $\mathbb{Z}$-algebra.

A more novel family of deformations comes from the following. Define $g_{ij} := f_i + e_j = e_i + f_j$.

Proposition B.9. *The Veronese subalgebra of Z_n corresponding to the sublattice spanned by $\langle g_{1,n+1}, \ldots, g_{n,n+1}\rangle$ is isomorphic to the multigraded coordinate ring of $(\mathbb{P}^1)^n$.*

Proof. Consider an element of degree

$$(r_1 + \cdots + r_n)h - (r_2 + \cdots + r_n)e_1 - \cdots - (r_1 + \cdots + r_{n-1})e_n \tag{B.12}$$

in Z_n. If $r_1 < 0$, then such an element vanishes at the first coordinate vector to higher order than its degree, and must therefore be 0. It follows that any nonzero element has $r_1, \cdots, r_n \geq 0$.

Let the generators of the coordinate ring of $\mathbb{P}^n$ be $y_1, \ldots, y_{n+1}$. Then we find more generally that an element of the above degree has degree at most r_i in y_i, and thus if we set y_{n+1} we obtain the inhomogeneous form of a multihomogeneous polynomial of multidegree $(r_1, \ldots, r_n)$. On the other hand, any monomial of degree $\leq r_i$ in y_i gives rise to an element of the desired space, and the result follows. □

Corollary B.10. *The subcategory of $\mathcal{A}$ with objects $\langle g_{1,n+1}, \ldots, g_{n,n+1}\rangle$ is a flat deformation of the natural $\mathbb{Z}^n$-algebra of $(\mathbb{P}^1)^n$.*

In fact, we can say more.

Proposition B.11. *The given flat deformation of $(\mathbb{P}^1)^n$ is generated in degrees $g_{1,n+1},\ldots,g_{n,n+1}$, and the relations form a quadratic Gröbner basis.*

Proof. We first note that the Gröbner basis property of $\mathcal{A}$ also holds relative to the ordering

$$[e_1]^{a_1}[f_1]^{b_1}\cdots[e_{n+1}]^{a_{n+1}}[f_{n+1}]^{b_{n+1}}; \tag{B.13}$$

here we use the expansion $[g_{i,j}] = [e_i][f_j] + [f_i][e_j]$ instead of $[g_{i,j}] = [e_i][f_j] + [e_j][f_i]$. Next, we observe that

$$[e_{n+1}][g_{i,n+1}] = [g_{i,n+1}][e_{n+1}] \tag{B.14}$$
$$[f_{n+1}][g_{i,n+1}] = [g_{i,n+1}][f_{n+1}]. \tag{B.15}$$

Indeed, using our standard basis for $[g_{i,n+1}]$, we need merely note that

$$[e_{n+1}]([e_i][f_{n+1}]) = [e_i][e_{n+1}][f_{n+1}] = ([e_i][f_{n+1}])[e_{n+1}] \tag{B.16}$$

and

$$[e_{n+1}]([f_i][e_{n+1}]) = ([e_{n+1}][f_i])[e_{n+1}] \subset [g_{i,n+1}][e_{n+1}]. \tag{B.17}$$

It is then easy to show by induction (using $[e_i][f_i] = [f_i][e_i]$) that

$$[g_{i,n+1}]^c = \bigoplus_{a+b=c} [e_i]^a[f_i]^b[e_{n+1}]^b[f_{n+1}]^a. \tag{B.18}$$

and then that

$$\begin{aligned}[g_{1,n+1}]^{c_1}\cdots[g_{n,n+1}]^{c_n} &= \bigoplus_{a_1+b_1=c_1} [e_1]^{a_1}[f_1]^{b_1}[e_{n+1}]^{b_1}[f_{n+1}]^{a_1}[g_{2,n+1}]^{c_2}\cdots[g_{n,n+1}]^{c_n} \\ &= \bigoplus_{a_1+b_1=c_1} [e_1]^{a_1}[f_1]^{b_1}[g_{2,n+1}]^{c_2}\cdots[g_{n,n+1}]^{c_n}[e_{n+1}]^{b_1}[f_{n+1}]^{a_1}\end{aligned} \tag{B.19}$$

But this implies by induction that

$$[c_1 g_{1,n+1} + \cdots + c_n g_{n,n+1}] = [g_{1,n+1}]^{c_1}\cdots[g_{n,n+1}]^{c_n}, \tag{B.20}$$

so that the algebra is generated as required.

We obtain n quadratic relations (starting from any given object) on these generators from the fact that $\dim[2g_{i,n+1}] = 3$ and $2n(n-1)$ quadratic relations from the fact that $\dim[g_{i,n+1}+g_{j,n+1}] = 4$, and these relations are clearly enough to reduce to the basis of ordered monomials. It follows that there are no more relations, and moreover that the given relations form a Gröbner basis. □

Remark. As we mentioned at the end of the previous section, the Gröbner basis property for these relations is a Yang-Baxter-type equation; the $n = 3$ case of the above can thus be viewed as giving a new proof of the 2×2 elliptic solution of the dynamical YBE [20].

Remark. The reader should be cautioned that the above constructions do not directly give non-commutative schemes without making some sort of choice of ample divisor. Although in the 2-dimensional case, the choice turned out to be mostly irrelevant, this is not the case for $n > 2$. The issue here is that the above $\mathbb{Z}^{n+2}$-algebra could be viewed both as an $n+1$-point blowup of $\mathbb{P}^n$ and as a 2-point blowup of $(\mathbb{P}^1)^n$, and these rational n-folds are no longer isomorphic for $n > 2$. For instance, the very ample line bundles $\mathcal{O}_X(d,d,d)$ on $(\mathbb{P}^1)^3$ have Euler characteristic $(d+1)^3$, while the corresponding line bundles $\mathcal{O}_X(3dh - 2de_1 - 2de_2 - 2de_3)$ on the blowup of $\mathbb{P}^3$ have the same global sections, but have Euler characteristic $\frac{d^3+6d^2+7d+2}{2}$.

We also obtain models for noncommutative rational surfaces as quotients of this $(\mathbb{P}^1)^n$. The nicest is probably the following. Again we work relative to an even blowdown structure for convenience.

We first consider deformations of (smooth) quintic del Pezzo surfaces.

Proposition B.12. *Let X_3 be a noncommutative rational surface such that no -2-curve is effective. Let $\mathcal{F}$ be the category obtained from $\mathcal{S}$ by insisting that all objects be integer linear combinations of s, f, $s+f-e_2-e_3$, $s+f-e_1-e_3$, $s+f-e_1-e_2$, and that all morphisms be* nonnegative *integer linear combinations of those elements. Then $\mathcal{F}$ is generated in those degrees, and is a quotient of a suitable $(\mathbb{P}^1)^5$ as constructed above.*

Proof. The claim about generation is essentially the only nontrivial claim, as it is easy to see that the morphisms satisfy the requisite relations. We need to show that any Hom space of degree in the relevant cone is generated as claimed. Since our generators are $W(E_4) \cong S_5$-invariant, we may restrict our attention to the intersection with the universal nef cone, which we may verify is the free monoid generated by the elements

$$f, s+f, 2s+2f-e_1-e_2, 3s+3f-2e_1-e_2-e_3, 4s+4f-2e_1-2e_2-2e_3. \tag{B.21}$$

(Here we should note that C_3 is a *nonnegative* combination of the generating degrees, but not an *integer* combination.) By our standard surjectivity results, it will suffice to show that Hom spaces of those five degrees are generated as required.

The first space *is* a generator, while $[s+f] = [s][f]$ is standard. That $[2s+2f-e_1-e_2] = [s+f][s+f-e_1-e_2]$ was discussed in the proof of Proposition B.1.

For the fourth space, we observe that

$$[3s+3f-2e_1-e_2-e_3] = [2s+2f-e_1-e_2][s+f-e_1-e_3]+[2s+2f-e_1-e_3][s+f-e_1-e_2], \tag{B.22}$$

since the right-hand side certainly surjects on leading coefficients, reducing to

$$[s+f-e_1] = [e_3][s+f-e_1-e_3] + [e_2][s+f-e_1-e_2]. \tag{B.23}$$

Since the four spaces are all in the given cone, the claim for this case follows from previous cases.

Finally, we claim that

$$[s+f-e_2-e_3][3s+3f-2e_1-e_2-e_3] = [4s+4f-2e_1-2e_2-2e_3]. \tag{B.24}$$

Again, this is surjective on leading coefficients, while

$$[s+f-e_2-e_3][s+f-e_1] = [2s+2f-e_1-e_2-e_3] \tag{B.25}$$

follows from the proof of Lemma 3.3, since none of $f-e_2-e_3$, $s-e_2-e_3$, e_1-e_2, e_1-e_3 are effective. □

Remark. Note that this monoid contains ample elements, so that we indeed obtain a model for the noncommutative quintic del Pezzo surface as a "closed subscheme" of a noncommutative $(\mathbb{P}^1)^5$.

There is also a construction for more general surfaces. We begin by constructing surjections from the full category $\mathcal{A}$. The functors will be straightforward to construct, but showing that the appropriate subcategory has the correct generators will require a technical lemma.

Lemma B.13. *Let $C \subset \mathbb{P}^1 \times \mathbb{P}^1$ be a smooth biquadratic curve, let $d, d' \in 1 + \mathbb{N}$, and suppose $p_1, \dots p_{d+d'}$ are distinct points of C. For $i \in \{1, \dots, d + d'\}$, let f_i be the nonzero bidegree $(1, 0)$ polynomial vanishing at p_i, and let g_i be the nonzero bidegree $(0, 1)$ polynomial vanishing at p_i. Then the polynomials $\prod_{i \in S} f_i \prod_{i \notin S} g_i$ with S ranging over d-element subsets of $\{1, \dots, d+d'\}$ span the space of bidegree (d, d') polynomials vanishing at $p_1, \dots, p_{d+d'}$.*

Proof. Suppose first that $d, d' \geq 2$. By induction, the claim holds for the point set $p_1, \dots, p_{d+d'-1}$ in bidegree $(d-1, d')$ and $(d, d'-1)$. Let $V_{d,d'}$ be the space of bidegree (d, d') polynomials vanishing at the base points, and let $V_{d-1,d'}$ and $V_{d,d'-1}$ be the corresponding spaces vanishing at the first $d+d'-1$ base points. We have $\dim V_{d,d'} = dd'+1$ as there are $(d-1)(d'-1)$ polynomials vanishing on C and a line bundle of degree $2(d + d')$ has $d + d'$ sections vanishing on the base points. We similarly have $\dim V_{d-1,d'} = (d-1)d' + 1$ and $\dim V_{d,d'-1} = d(d'-1) + 1$.

Now, we want to show that the map

$$V_{d-1,d'} \oplus V_{d,d'-1} \xrightarrow{(f_{d+d'}, g_{d+d'})} V_{d,d'} \tag{B.26}$$

is surjective, and given what we have shown about the dimensions, it would suffice to show that the kernel has dimension $(d-1)(d'-1)$. Now, the kernel consists of bidegree $(d-1, d'-1)$ polynomials h such that $f_{d+d'}h$ and $g_{d+d'}h$ both vanish at the base points. Since the only common zero of $f_{d+d'}$ and $g_{d+d'}$ is $p_{d+d'}$, we find that W is the space of polynomials of bidegree $(d-1, d'-1)$ vanishing at the first $d + d' - 1$ base points, and we may calculate as before that $\dim W = (d-1)(d'-1)$ as required.

It remains to consider the case $d = 1$. The case $d' = 1$ is straightforward (there are only three distinct PGL_2^2 orbits to consider), so we may assume $d' > 1$. Then we still have $\dim V_{1,d'} = d' + 1$, $\dim V_{1,d'-1} = d'$, and multiplication by $g_{d'+1}$ injects $V_{1,d'-1}$ in $V_{1,d'}$. It will thus suffice to show that $h = f_{d'+1} \prod_{1 \leq i \leq d'} g_i$ is not in the image of $V_{1,d'-1}$. If $g_i \neq g_{d'+1}$ for $1 \leq i \leq d$, then h is not even a multiple of $g_{d'+1}$, so the claim is immediate. Otherwise, suppose $g_{d'} = g_{d'+1}$. But then we see that none of the factors of

$$f_{d'+1} \prod_{1 \leq i < d'} g_i \tag{B.27}$$

vanish at $p_{d'}$, so again this polynomial is not in $V_{1,d'-1}$. □

Proposition B.14. *Let $X_m = X_{\rho;q;C}$ be a noncommutative surface equipped with an even blowdown structure such that no root of the forms $s - f$, $f - e_i - e_j$, $s - e_i - e_j$, $e_i - e_j$ is effective. Let $\mathcal{F}$ be the subcategory associated to the monoid generated by $f - e_1, \dots, f - e_m$ and $s - e_1, \dots, s - e_m$. Then there is a functor from a special case of $\mathcal{A}$ to $\mathcal{F}$ acting on objects as the linear transformation*

$$\begin{aligned} h &\mapsto s + (m-1)f - e_1 - \cdots - e_m \\ e_i &\mapsto f - e_i, \end{aligned}$$

and this functor is surjective on Hom *spaces.*

Proof. As above, the existence of a functor is straightforward, and the nontrivial part of the proof is showing that $\mathcal{F}$ is generated by elements of the given degrees. The relevant monoid consists of elements

$$ds + d'f - \sum_i r_i e_i \tag{B.28}$$

where $d, d', r_1, \dots, r_m \geq 0$ and $d + d' = \sum_i r_i$, so we need to prove that every Hom space of that degree is spanned as required.

Suppose first that $\max(r_i) = 1$. Then without loss of generality we have an element of the form

$$ds + d'f - e_1 - \cdots - e_m \tag{B.29}$$

with $d + d' = m$ (as otherwise we can act by S_m and then blow down) and $d \leq d'$ (as otherwise we can reflect in $s - f$). Since $mf - e_1 - \cdots - e_m$ has negative intersection with each -1-curve of the form $f - e_i$, we have $[mf - e_1 - \cdots - e_m] = [f - e_1] \cdots [f - e_m]$ as required, giving the case $d = 0$.

Otherwise, for any d-element subset $S \subset \{1, \ldots, m\}$, we have a 1-dimensional space

$$\prod_{i \notin S} [f - e_i] \prod_{i \in S} [s - e_i] \subset [ds + d'f - e_1 - \cdots - e_m] \tag{B.30}$$

If we left-multiply by $\prod_i [e_i]$, we obtain the space

$$\prod_{i \notin S} ([e_i][f - e_i]) \prod_{i \in S} ([e_i][s - e_i]) \subset [df][d's] \tag{B.31}$$

It follows from the Lemma that as S varies, these span $ds + d'f - e_1 - \cdots - e_m$.

It remains to consider the case $\max(r_i) > 1$, say $r_1 > 1$. In that case, we wish to show

$$\begin{aligned} [ds + d'f - r_1e_1 - \cdots - r_me_m] = \quad & [f - e_1][ds + (d' - 1)f - (r_1 - 1)e_1 - \sum_{2 \leq i \leq m} r_ie_i] \\ & + [ds + (d' - 1)f - (r_1 - 1)e_1 - \sum_{2 \leq i \leq m} r_ie_i][s - e_1] \end{aligned} \tag{B.32}$$

Since $s - f$ is ineffective, this holds for leading coefficients, so reduces to showing

$$\begin{aligned} & [(d-2)s + (d'-3)f - (r_1 - 1)e_1 - \cdots - (r_m - 1)e_m] \\ = \quad & [f - e_1][(d-2)s + (d'-3)f - (r_1 - 2)e_1 - \sum_{2 \leq i \leq m} (r_i - 1)e_i] \\ & + [(d-2)s + (d'-3)f - (r_1 - 2)e_1 - \sum_{2 \leq i \leq m} (r_i - 1)e_i][s - e_1], \end{aligned} \tag{B.33}$$

which holds by induction. □

Corollary B.15. *Let X_m be a blowup of noncommutative $\mathbb{P}^2$ in $m+1$ points in distinct $q^{\mathbb{Z}}$-orbits, and let $\mathcal{F}$ be the subcategory associated to the monoid generated by $h - e_1, \ldots, h - e_{m+1}$. Then $\mathcal{F}$ is generated in those degrees, and is a quotient of the above noncommutative $(\mathbb{P}^1)^{m+1}$.*

Proof. If we blow up two more points in sufficiently general position, then modify the blowdown structure appropriately, we may rephrase this in terms of the monoid generated by $s + f - e_1 - e_2, \ldots, s + f - e_1 - e_{m+2}$ on X_{m+2}. The resulting surface satisfies the hypotheses of Proposition B.14, and restricting the functor to the corresponding $(\mathbb{P}^1)^{m+1}$ (which is independent of the auxiliary blowups) gives the desired result. □

Remark. Up to the action of $W(E_{m+1})$, this case and the quintic del Pezzo case are the only configurations of rational pencils such that any two of the pencils have intersection 1. Indeed, there is certainly an even blowdown structure in which the first pencil is f, at which point the remaining pencils have the form $s + d'f - \sum_i r_ie_i$. The action of D_m takes the next pencil to s, at which point the remaining pencils must all have the form $s + f - e_i - e_j$ for $i < j$. We thus obtain a graph on the m points such that any two edges share a vertex. The only two cases are the complete graph K_3, corresponding to the quintic del Pezzo surface case, and star graphs, which relative to a $\mathbb{P}^2$ basis have the above form.

References

[1] A. B. Altman and S. L. Kleiman. Compactifying the Picard scheme. *Adv. in Math.*, 35(1):50–112, 1980.

[2] A. Andreotti and T. Frankel. The Lefschetz theorem on hyperplane sections. *Ann. of Math. (2)*, 69:713–717, 1959.

[3] M. Artin, J. Tate, and M. Van den Bergh. Some algebras associated to automorphisms of elliptic curves. In *The Grothendieck Festschrift, Vol. I*, volume 86 of *Progr. Math.*, pages 33–85. Birkhäuser Boston, Boston, MA, 1990.

[4] M. Artin, J. Tate, and M. Van den Bergh. Modules over regular algebras of dimension 3. *Invent. Math.*, 106(2):335–388, 1991.

[5] M. Artin and J. J. Zhang. Noncommutative projective schemes. *Adv. Math.*, 109(2):228–287, 1994.

[6] M. Artin and J. J. Zhang. Abstract Hilbert schemes. *Algebr. Represent. Theory*, 4(4):305–394, 2001.

[7] M. P. Bellon and C.-M. Viallet. Algebraic entropy. *Comm. Math. Phys.*, 204(2):425–437, 1999.

[8] M. Bhargava and W. Ho. Coregular spaces and genus one curves. *Camb. J. Math.*, 4(1):1–119, 2016.

[9] G. D. Birkhoff. The generalized Riemann problem for linear differential equations and the allied problems for linear difference and q-difference equations. *Proc. Am. Acad. Arts Sci.*, 49(9):521–568, 1913.

[10] A. I. Bondal and M. M. Kapranov. Representable functors, Serre functors, and reconstructions. *Math. USSR-Izv.*, 35(3):519–541, 1990.

[11] A. I. Bondal and A. E. Polishchuk. Homological properties of associative algebras: the method of helices. *Izv. Ross. Akad. Nauk Ser. Mat.*, 57(2):3–50, 1993.

[12] F. Bottacin. Poisson structures on moduli spaces of sheaves over Poisson surfaces. *Invent. Math.*, 121(2):421–436, 1995.

[13] T. Bridgeland, A. King, and M. Reid. The McKay correspondence as an equivalence of derived categories. *J. Amer. Math. Soc.*, 14(3):535–554 (electronic), 2001.

[14] D. Chan and A. Nyman. Non-commutative Mori contractions and $\mathbb{P}^1$-bundles. *Adv. Math.*, 245:327–381, 2013.

[15] S. Cooper. The Askey-Wilson operator and the ${}_6\Phi_5$ summation formula. *South East Asian J. Math. Math. Sci.*, 1(1):71–82, 2002.

[16] J. F. van Diejen. Integrability of difference Calogero-Moser systems. *J. Math. Phys.*, 35(6):2983–3004, 1994.

[17] P. Etingof and C.-H. Eu. Koszulity and the Hilbert series of preprojective algebras. *Math. Res. Lett.*, 14(4):589–596, 2007.

[18] P. Etingof, W. L. Gan, and A. Oblomkov. Generalized double affine Hecke algebras of higher rank. *J. Reine Angew. Math.*, 600:177–201, 2006.

[19] P. I. Etingof. Galois groups and connection matrices of q-difference equations. *Electron. Res. Announc. Amer. Math. Soc.*, 1(1):1–9 (electronic), 1995.

[20] G. Felder. Elliptic quantum groups. In *XIth International Congress of Mathematical Physics (Paris, 1994)*, pages 211–218. Int. Press, Cambridge, MA, 1995.

[21] G. Felder and A. Varchenko. The elliptic gamma function and $\mathrm{SL}(3, \mathbf{Z}) \ltimes \mathbf{Z}^3$. *Adv. Math.*, 156(1):44–76, 2000.

[22] T. Graber and E. M. Rains. Perverse Hilbert schemes of elliptic surfaces. in preparation.

[23] B. Harbourne. Anticanonical rational surfaces. *Trans. Amer. Math. Soc.*, 349(3):1191–1208, 1997.

[24] R. Hartshorne. *Algebraic geometry*. Springer-Verlag, New York-Heidelberg, 1977. Graduate Texts in Mathematics, No. 52.

[25] J. C. Hurtubise and E. Markman. Elliptic Sklyanin integrable systems for arbitrary reductive groups. *Adv. Theor. Math. Phys.*, 6(5):873–978 (2003), 2002.

[26] D. Huybrechts and M. Lehn. *The geometry of moduli spaces of sheaves*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, second edition, 2010.

[27] A. Ishii and H. Uehara. Autoequivalences of derived categories on the minimal resolutions of A_n-singularities on surfaces. *J. Differential Geom.*, 71(3):385–435, 2005.

[28] M. E. H. Ismail, E. M. Rains, and D. Stanton. Orthogonality of very well-poised series. In preparation.

[29] K. Kajiwara, T. Masuda, M. Noumi, Y. Ohta, and Y. Yamada. Point configurations, Cremona transformations and the elliptic difference Painlevé equation. In *Théories asymptotiques et équations de Painlevé*, volume 14 of *Sémin. Congr.*, pages 169–198. Soc. Math. France, Paris, 2006.

[30] I. M. Krichever. Analytic theory of difference equations with rational and elliptic coefficients and the Riemann-Hilbert problem. *Uspekhi Mat. Nauk*, 59(6(360)):111–150, 2004.

[31] S. G. Langton. Valuative criteria for families of vector bundles on algebraic varieties. *Ann. of Math. (2)*, 101:88–110, 1975.

[32] M. Lehn. On the cotangent sheaf of Quot-schemes. *Internat. J. Math.*, 9(4):513–522, 1998.

[33] W. Lowen and M. Van den Bergh. Deformation theory of abelian categories. *Trans. Amer. Math. Soc.*, 358(12):5441–5483 (electronic), 2006.

[34] V. B. Mehta and A. Ramanathan. An analogue of Langton's theorem on valuative criteria for vector bundles. *Proc. Roy. Soc. Edinburgh Sect. A*, 96(1-2):39–45, 1984.

[35] D. Mumford. *Lectures on curves on an algebraic surface*. Annals of Mathematics Studies, No. 59. Princeton University Press, Princeton, N.J., 1966.

[36] T. A. Nevins and J. T. Stafford. Sklyanin algebras and Hilbert schemes of points. *Adv. Math.*, 210(2):405–478, 2007.

[37] A. Nyman. Serre duality for non-commutative $\mathbb{P}^1$-bundles. *Trans. Amer. Math. Soc.*, 357(4):1349–1416 (electronic), 2005.

[38] A. V. Odesskiĭ. Elliptic algebras. *Russian Math. Surveys*, 57(6):1127–1162, 2002.

[39] A. Okounkov and E. Rains. Noncommutative geometry and Painlevé equations. *Algebra Number Theory*, 9(6):1363–1400, 2015.

[40] D. Orlov. Smooth and proper noncommutative schemes and gluing of dg categories. arXiv:1402.7364v2.

[41] C. Ormerod and E. M. Rains. An elliptic Garnier system. arXiv:1607.07831.

[42] C. Praagman. Fundamental solutions for meromorphic linear difference equations in the complex plane, and related problems. *J. Reine Angew. Math.*, 369:101–109, 1986.

[43] E. Rains and S. Ruijsenaars. Difference operators of Sklyanin and van Diejen type. *Comm. Math. Phys.*, 320(3):851–889, 2013.

[44] E. M. Rains. The birational geometry of noncommutative ruled surfaces. In preparation.

[45] E. M. Rains. Birational morphisms and Poisson moduli spaces. arXiv:1307.4032.

[46] E. M. Rains. Elliptic double affine Hecke algebras and their spherical algebras. In preparation.

[47] E. M. Rains. Generalized Hitchin systems on rational surfaces. arXiv:1307.4033.

[48] E. M. Rains. Multivariate quadratic transformations and the interpolation kernel. arXiv:1408.0305.

[49] E. M. Rains. Transformations of elliptic hypergeometric integrals. *Ann. of Math. (2)*, 171(1):169–243, 2010.

[50] E. M. Rains. An isomonodromy interpretation of the hypergeometric solution of the elliptic Painlevé equation (and generalizations). *Symmetry, Integrability, and Geometry: Methods and Applications*, 7(088):24 pages, 2011. arXiv:0807.0258.

[51] E. M. Rains and V. P. Spiridonov. Determinants of elliptic hypergeometric integrals. *Funktsional. Anal. i Prilozhen.*, 43(4):67–86, 2009.

[52] S. N. M. Ruijsenaars. First order analytic difference equations and integrable quantum systems. *J. Math. Phys.*, 38:1069–1146, 1997.

[53] S. N. M. Ruijsenaars. Hilbert-Schmidt operators vs. integrable systems of elliptic Calogero-Moser type IV. The relativistic Heun (van Diejen) case. *SIGMA Symmetry Integrability Geom. Methods Appl.*, 11:Paper 004, 78, 2015.

[54] H. Sakai. Rational surfaces associated with affine root systems and geometry of the Painlevé equations. *Comm. Math. Phys.*, 220(1):165–229, 2001.

[55] M. J. Schlosser and M. Yoo. Elliptic hypergeometric summations by Taylor series expansion and interpolation. arXiv:1602.09027.

[56] S. Schröer. Topological methods for complex-analytic Brauer groups. *Topology*, 44(5):875–894, 2005.

[57] P. Seidel and R. Thomas. Braid group actions on derived categories of coherent sheaves. *Duke Math. J.*, 108(1):37–108, 2001.

[58] E. K. Sklyanin. Some algebraic structures connected with the Yang-Baxter equation. *Funktsional. Anal. i Prilozhen.*, 16(4):27–34, 96, 1982.

[59] E. K. Sklyanin. Some algebraic structures connected with the Yang-Baxter equation. Representations of a quantum algebra. *Funktsional. Anal. i Prilozhen.*, 17(4):34–48, 1983.

[60] V. P. Spiridonov. On the elliptic beta function. *Russian Math. Surveys*, 56(1):185–186, 2001.

[61] V. P. Spiridonov. The continuous biorthogonality of an elliptic hypergeometric function. *St. Petersburg Math. J.*, 20(5):791–812, 2009.

[62] V. P. Spiridonov and S. O. Warnaar. Inversions of integral operators and elliptic beta integrals on root systems. *Adv. Math.*, 207(1):91–132, 2006.

[63] A. N. Tyurin. Symplectic structures on the moduli spaces of vector bundles on algebraic surfaces with $p_g > 0$. *Izv. Akad. Nauk SSSR Ser. Mat.*, 52(4):813–852, 896, 1988.

[64] M. Van den Bergh. Blowing up of non-commutative smooth surfaces. math.QA/9809116v1.

[65] M. Van den Bergh. Non-commutative $\mathbb{P}^1$-bundles over commutative schemes. *Trans. Amer. Math. Soc.*, 364(12):6279–6313, 2012.

[66] J.-L. Verdier. Des catégories dérivées des catégories abéliennes. *Astérisque*, 239:xii+253 pp., 1996. (A slightly edited reprint of Verdier's 1967 thesis).

[67] Y. Yamada. A Lax formalism for the elliptic difference Painlevé equation. *SIGMA Symmetry Integrability Geom. Methods Appl.*, 5:Paper 042, 15, 2009.